Proto Balance Mail – SMTP/POP Load Balancer

http://www.protonet.co.za/

Mail User Manual

Proto Balance Mail - User Manual

March 28, 2010

Contents

1 Introduction

1.1 Overview

Proto Balance Mail is an email load balancing solution for companies, universities and ISPs that host between 1,000 and 1,000,000 mail accounts.

Proto Balance Mail is unique in offering scalability, flexibility, fall-over safety, and security using your existing hardware and network infrastructure. Other mail solutions work with expensive SAN arrays and clustered file-systems whereas Proto Balance Mail can provide better performance with entry-level commodity hardware at a fraction of the cost.

Proto Balance Mail is protected by patents filed in the U.S.A. and the European Union.

1.2 Installation and start-up

To install Proto Balance Mail, download an installer for your operating system from the Free Trial tab. On Windows you can run the installer directly. On Unix systems (including MacOS X) you should give the installer executable permissions and install it as root. Type the following command at the command-prompt:

```
chmod a+x install.bin
./install.bin
```

Then follow the instructions.

To start up Proto Balance Mail, simply type,

```
protobalance
```

You can now point your web browser to `http://127.0.0.1:8080/` to view Proto Balance Mail's configuration and diagnostic interface.

See the generic user manual for Proto Balance to get information on running Proto Balance Mail with various settings for high-performance.

2 How Proto Balance Mail Works

Proto Balance Mail distributes email load over a number of servers such that the email of a particular user is always delivered to, and can always be retrieved (POPped) from, the *same* server.

This means if you host a total of 100,000 mail accounts at your organization, you can host 10,000 mail accounts on each of 10 servers and Proto Balance Mail will ensure that SMTP connections are routed to the correct server. Proto Balance will also ensure that POP connections are routed to the server on which the POP account is hosted.

Proto Balance Mail does not store or queue email, but works by manipulating the SMTP session according to a patented method. No other load balancing product can do this.

Proto Balance Mail can deliver hundreds of emails *per second* and scale to a large number of back-end servers. Multiple Proto Balance Mail instances can exist on the same segment - these will automatically detect each other and keep their data in using replication.

Each Proto Balance Mail instance can handle up to 10,000 concurrent incoming SMTP and POP connections.

3 More About How Proto Balance Mail Works

With Proto Balance Mail, your own email servers receive and store email as shown in the deployment view below. To give an example of its operation, consider the situation of a Google-Mail user sending email to you. Let's say that the user is harryp@gmail.com and that your email account is headmaster@mydomain.com.

The user harryp@gmail.com will compose and send the email; Google's email server will queue the email. Google's email server will see that the recipient address is headmaster@mydomain.com and will hence lookup the MX record for mydomain.com. The MX record for mydomain.com will be the machine on which you have installed Proto Balance Mail. Google's email server will connect to Proto Balance Mail and begin an email session to transmit the email. Proto Balance Mail will see that you have previously configured an email account headmaster@mydomain.com. The email account will have a back-end email server machine associated to it. The back-end mail-server can be Microsoft Exchange, PostFix, Sendmail, QMail, or any other mail server software you have previously installed. Proto Balance Mail will transparently forward the email session to your back-end mail server. Your back-end mail server will append the email to your mailbox.

At your PC or Mac you will then try to retrieve your new email, aka your "Inbox". Your email client will have a user name and password configured to retrieving email using, say, POP3. Your email client will have configured for "incoming email" the server where Proto Balance Mail is running; and will connect to this server. Proto Balance Mail will authenticate you and transparently direct your email retrieval session to the same back-end email server where your mail box is stored.

Your email client will display the email to you.

You may have thousands or millions of email accounts, but each email server need only host a portion of them. The email servers need not know about each other - they only need to know about the email accounts they alone are responsible for.

In this way, you can divide-up your email accounts into groups small enough to be managed efficiently by *one* email server. A typical email server works well for up to 1000 email accounts. Most system administrators notice that when the number of email accounts grows to larger than 1000, then significant performance degradation occurs - a desktop user POPing mail becomes very sluggish, the email server becomes overloaded, the hard drive starts to thrash, and the machine becomes unstable. Most importantly, it becomes extremely difficult to move or upgrade the server because so much email is stored within it.

With Proto Balance Mail you can add as *many* email servers as you like, and the email servers seamlessly operate as a cluster. In particular

* You do not need NFS or other remote file-systems.

* You do not need a clustered or other shared file-system.

* You do not need a Storage Area Network (SAN); expensive - ouch!

* You do not need to run any specialized software on your mail server: just your regular email handling software will do.

The only administrative action required is to maintain Proto Balance Mail's built-in database of email accounts, along with the email server on which each account is hosted.

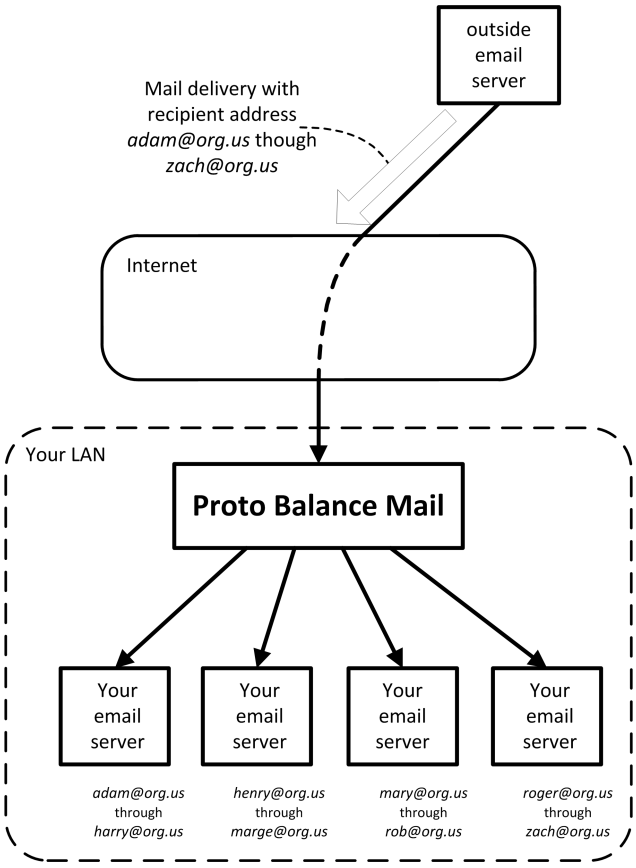

4 First-time Configuration

To configure Proto Balance Mail for incoming email takes only a few minutes.

After installing and invoking Proto Balance Mail you should be able to login to the configuration web page. See the main user manual for more information on installation and startup.

Note that on Unix, Proto Balance Mail requires more semaphores than other Proto Balance Products.

You need to start Proto Balance Mail with traffic logging enabled. Make a new folder/directory "C:\maillog" (Windows) or "/var/log/maillog" (Unix) use the command-line option:

-monthly C:\maillog

or

-monthly /var/log/maillog

4.1 Step 1: Create an SMTP cluster

Proto Balance Mail needs to listen on port 25 for incoming mail delivery attempts. Therefore your first action should be to click on the "Add SMTP Cluster" tab, which adds a listener for port 25. In the form, only the fields with a yellow background are mandatory. You can leave all other fields at their default values.

The field "SMTP host name" is the official name of the mail server used in status messages on the wire. This should be a fully qualified, DNS-resolvable host name. "Unique Identifier" can be anything.

It is important that Proto Balance Mail does not attempt to deliver mail messages larger than your existing email servers can accept, therefore you must set "Max mail size in kilobytes" to be smaller than the maximum email message size allowed by your mail server.

For testing purposes, you should turn spam filtering off. Set "Default spam filtering" to "Disabled".

4.2 Step 2: Add a back-end mail server

You now need to add at least one back-end mail server that will accept email deliveries. This mail server needs to be configured with user accounts and mail boxes and be ready to accept email deliveries. If you have an existing mail server then this should work fine.

Click on "Clusters" and then on "Add box".

Fill in the IP address of your email server next to "TCP connect address".

Fill in the listening port of your email server next to "TCP connect port" - this will usually be port 25.

The field "Unique Identifier" is important. This will be a text string used to identify the box to which an email is sent for a particular recipient. You should choose something short and descriptive. This will be discussed in the next step.

After adding the box, click on the green traffic light to enable it.

4.3 Step 3: Adding an email account

Click on "Email Database" and scroll down to the "Add Email" form. Within this form you can add a new user account.

Next to "Email address" you need to put the full email address of the user account. As an example, let us say that this user is "henry@org.us".

Next to "Box Name" you need to insert the unique identifier of the previous step. This will mean that incoming mail for this user will be delivered to that mail server.

The "Password" field is for POP authentication and outgoing email. This should be set to the same as the password for the same user account on your mail server. Outgoing mail and POP will be discussed later.

After you have added the email account, you can click on "Search Email" to view or modify the account.

4.4 Step 4: Testing mail delivery

You can now test email delivery by switching your MX record for "org.us" to Proto Balance Mail, and then attempting to send an email to "henry@org.us". Alternatively, you can use telnet to connect to port 25 of Proto Balance Mail, as follows:

220 smtp.@org.us ESMTP
HELO
250 ok
MAIL FROM:<admin@org.us>
250 ok
RCPT TO:<henry@org.us>
250 ok
DATA
354 Start mail input; end with <CRLF>.<CRLF>
Subject: test


```
test
.
250 [MAIL1:250_ok_1262775853_qp_13377__]
quit
221 bye
```

Under "mylogs" your log file will then contain a line:

```
2010-01-06 Wed 10:32:48.365: 10.11.12.13:
        admin@org.us => henry@org.us
        "test":0kB success - 250 [MAIL1:250_ok_1262775853_qp_13377__]
```

As you can see, Proto Balance Mail did not queue the mail itself, but directed the SMTP connection directly the appropriate mail server.

5 Bulk Importing of Thousands of Email accounts

After performing the above steps you can look under your configuration directory (see the -directory command-line setting) and find the file email.csv.

The file email.csv is a tab-delimited file containing all email accounts added through the web page. It is a "flat database file". Other software may have database files that are in binary format and are not human readable, but this file *is* human readable.

You can flush and rewrite email.csv by clicking on "Vacuum Databases" under the "Email Database" tab.

You can view and modify email.csv so long as you keep to the precise format and provided that you vacuum the database before you make modifications. Click on "Re-read Database" when you are done with modifications. Note that "Re-read Database" deletes all records before importing email.csv. Therefore there may be a short period when email may fail. You should therefore use "Re-read Database" only during maintenance cycles or during non-critical periods.

To add a large number of email accounts you need to export them as a tab-delimited file. The file should have three columns seperated by ASCII character 9 decimal. The columns are the email address, the password, and the unique identifier of the box from Step 2 above.

If you have your email accounts stored in a database, you need to consult your database documentation to allow you to export your database as plain text.

An example may be:

```
henry@org.us jdfu2bnc MAIL1
bob@dutchcorp.nl hgew352i MAIL3
alice@org.us pjdyw4f1 MAIL3
trudy@org.us zdqo2pr9 MAIL2
roger@example.com asdf345w MAIL2
percy@somewhere.us 768fds2i MAIL1
```

Extra whitespace anywhere will be interpreted literally, therefore be sure that there are no spurious ASCII space characters anywhere in the file.

Invalid lines will be skipped.

Once you have your exported file in the proper format, you can overwrite email.csv and then click on "Re-read Database". You can find your newly-important email accounts by clicking on "Search Email".

6 Valid Email Addresses

Proto Balance Mail has a strict definition of what is considered to be a valid email address. Email addresses that do not fit this definition may not be entered as an email account, or appear as a sender or recipient.

Email addresses are case insensitive.

The domain part of the email address (the part after the @ symbol) may contain any of the characters

0 through 9
a through z
. _ -

The local part of the email address (the part before the @ symbol) may contain any characters except

@ ( ) [ ] \ ; : , < > " ’

The local part may also not contain characters from 0 through 32 decimal and from 127 through 159 decimal. The local part may also not contain adjacent punctuation characters.

7 Configuring POP

To configure Proto Balance Mail so that your clients can retrieve their email (POP), proceed with the following steps.

7.1 Step 1: Create an POP3 cluster

Click on "Add POP3 Cluster" to create a new listener on the POP3 port. The form values are similar to that for the SMTP cluster.

7.2 Step 2: Add a back-end mail server

Under the "Clusters" tab, click on "Add box" next to your POP cluster.

The name of the box must be entered next to "Unique Identifier". This must be identical to the same box as you entered in "Step 2" for your SMTP cluster so that POP connections can be directed to the same back-end email server for the same recipient.

7.3 Step 3: Configure your email client

You can now attempt to POP email with your email client. Using the examples earlier, the mail account "henry@org.us" with password "jdfu2bnc" is available.

You can also test the POP account be issuing a telnet to Proto Balance Mail to port 110, as follows:

+OK <1262856808170@mail.localdomain>
user henry@org.us
+OK
pass jdfu2bnc
+OK
quit
+OK

The "+OK" in response to "pass" means that Proto Balance Mail successfully authenticated the user "henry@org.us" using the password previously entered into the email database. It also means that Proto Balance Mail was able to connect to the back-end mail server and login using henry@org.us and the same password "jdfu2bnc".

Note that the password in the email database *must* exactly match the password of the same email account known to your mail server.

Proto Balance Mail supports the POP3 authentication mechanisms of "APOP", "LOGIN", and "USER". However, Proto Balance Mail will itself connect to the back-end mail server using "USER" authentication. Hence your back-end mail server *must* support "USER" authentication. To very that it does, issue the same telnet session above directly to your back-end mail server.

8 Configuring Outgoing SMTP

Outgoing SMTP configuration is not critical to the handling of mail because it does not involve a mailbox. However the correct configuration of outgoing SMTP has certain advantages:

1. Whitelists between senders and recipients are created such that reply emails do not incur delays. In other words, if henry@org.us sends email through Proto Balance Mail to mary@gmail.com, then Proto Balance Mail will record the association between these two persons in its internal database. Thereafter, mary@gmail.com can send email to henry@org.us and be sure that her email will never be flagged as SPAM - it will be delivered immediately.

2. Proto Balance Mail can be configured to select a random outgoing email server from among the back-end servers flagged with "SMTP relay" - [Enabled]. This will load balance outgoing deliveries over all such servers.

3. Alternatively, Proto Balance Mail can be configured to use the same outgoing email server as where the mail-box is stored for that sender. Hence, if "henry@org.us" is sending email then Proto Balance Mail will direct the session to the same server as where his mailbox is stored. For this to work, in your cluster configuration, you must set "How to choose relay box" to [Use box in per-email config]. This setting has the advantage that problems sending mail can be easily diagnosed - it is easy to determine which box handled the outgoing delivery.

Your back-end email servers *must* be configured to trust Proto Balance Mail as follows:

1. Write down the IP address of the machine on which Proto Balance Mail is running.

2. Find the configuration settings for your email software on your back-end mail servers.

3. Find the setting for IP address ranges that are allowed to use your back-end mail servers to relay mail.

4. Add the IP address to this range.

You can configure an email client to use Proto Balance Mail for outgoing email delivery. This is similar to POP authentication in the above section. Proto Balance Mail supports the authentication mechanisms of "LOGIN", "PLAIN", and "CRAM-MD5".

Once all your email clients are using Proto Balance Mail for outgoing SMTP, it is recommended that, except for Proto Balance Mail itself, you completely block port 25 access to your back-end mail servers. This means that all email passes through Proto Balance Mail *only*.

(As a security precaution, one ISP did in fact block *all* port 25 *traffic* from it's DSL clients. This was done in order to protect it's dynamic address ranges from being black-listed due to Mail-bots. The only allowed port 25 traffic was through Proto Balance Mail.)

9 Multiple Proto Balance Mail Machines

To have complete fall-over safety with "no-single-point-of-failure", you can deploy multiple Proto Balance Mail machines that manage the same cluster of back-end mail servers. This means that you can have an MX record for each.

If you install more than one Proto Balance Mail machine on the same LAN (the same segment of your network) then Proto Balance Mail will automatically find other instances using broadcast discovery packets. Such discovered machines are called *sisters* and are visible under the Info tab next to the "Sisters" row of the top info box.

Proto Balance Mail will keep the databases of all sister machines synchronized. Specifically, the following data sets are kept replicated to and from all sister machines:

1. Email account database. 2. Email alias database. 3. Greylist records. 4. Whitelist records. 5. IP access records for spam analysis. 6. Account access records.

This means that if you add or modify a new email account to one Proto Balance Mail instance, then *all* sisters will instantly be updated with the identical account.

A machine on a different segment will *not* receive broadcast packets and will therefore not automatically be discovered. Such machines can be entered manually into the setting "Sister list" under the Info tab. For example:

> Sister list (...): 23.24.25.26, 69.68.67.66, 12.13.14.15

Next to "Sisters" in the top info box you can see the list of discovered sister machines. An example display is as follows:

> 12.13.14.15(tx43516,rx92675,q13,a3244)

This means that sister 12.13.14.15 is connected and we have sent to it 43516 transaction records, received 92675 transaction records, and currently have 13 records queued for sending and allocated 3244 bytes.

This outgoing queue is important: should sister 12.13.14.15 become unavailable then the queue will grow very large. When the queue reaches 50 megabytes in size, Proto Balance Mail will assume the sister is permanently discontinued.

The benefit of the transaction queue is that you can take down a sister and perform maintenance; and when the sister is restored all queued records will be flushed and bring the sister's databases up to a current state. This feature will work provided that the queue does not reach 50,000,000 bytes.

10 Alias Addresses

All organizations require alias addresses. For instance, if someone sends email to the address "accounts@org.us", then both "brad@@org.us" and "janet@org.us" should receive it. This might be for reasons of processes within your organization.

Configuring alias addresses can be done by going to the [Email Database] tab; filling in a new alias expansion; and clicking on the [Add Alias] button. The "Alias list" can be a comma-separated list.

The "Alias list" should be short. I.e. you should not try implement a large mailing list using this feature. Instead use dedicate mailing list software if you have this requirement.

If you need to bulk-add aliases (say from your previous mail handling software), first read the section "Bulk Importing of Thousands of Email accounts" above. Be sure to "Vacuum" and "Re-read" your database as required.

The flat database file used for alias information is "alias.csv" under the directory specified by the -directory command-line setting.

The file has a two-column format such as:

accounts@org.us brad@@org.us,janet@org.us
abuse@example.com admin@example.com
info@org.us john@org.us,peter@gmail.com
everyone@org.us info@org.us,accounts@org.us

Use the ASCII character 9 decimal to separate the two columns. Aside from newline characters at the end of each row, there should be no other white-space characters in the file.

Note that you can freely associate aliases to other aliases in a recursive fashion. Loops of aliases are handled safely.

11 Proto Balance Mail as a Windows Service

Proto Balance Mail installs as a Windows Service as well as a regular program under your Start menu. This means you can start Proto Balance Mail by going through the menus:

[Start] - [Control Panel] - [Administrative Tools] - [Services] - PROTOBALANCE

If you right click on PROTOBALANCE and select [Properties] - [Recovery], you can set Proto Balance Mail to restart on failure.

12 Proto Balance Mail Command-Line as a Windows Service

A windows service does not have a proper command-line. To set the command-line arguments, edit the file CMDLINE under the same directory as your executable. This is usually C:\Program Files\Proto Balance\CMDLINE

With the file CMDLINE you can enter options for the main log file, and mail transaction log files.

13 Proto Balance Mail Watchdog Timer

Under Unix systems you may want to enable a watchdog timer for cases where Proto Balance Mail might exit for reasons unknown. This would normally only happen for new and unstable releases.

The command-line option -watchdog will start an additional parent process to monitor and restart Proto Balance Mail if it should exit for any reason. If you enable this option you can never kill Proto Balance Mail except by killing both the parent and child processes.

14 Log Files

The usual log file command-line argument -logfile can be used to specify the primary log file of Proto Balance Mail. This log file contains low-level error information and is not especially voluminous.

An additional log file, specifically to show email transactions, is useful for diagnostics, statistics, performance measurement, and forensics. This is specified with the command-line argument:

-monthly /var/log/maillog

Or on Windows:

-monthly C:\maillog

You should be sure to create this directory beforehand.

These log files can grow to a large size. For files that rotate daily, use instead the command-line argument,

-daily /var/log/maillog

Or on Windows:

-daily C:\maillog

This will create a new log file every day at 00h00 GMT.

The log files will show all important protocol interactions and their client connection address. It is advisable to study these files to see how your email is performing.

15 Settings Reference

All Proto Balance Mail configuration settings are explained in this section.

15.1 Cluster configuration settings

Unique Identifier: This is the cluster name. It merely needs to be unique for purposes of identification; many XML commands will use this identifier.

SMTP host name: Used for SMTP responses. The SMTP protocol dictates that certain on-the-wire interactions include a host name, and that this host name is the fully qualified domain name (FQDN) of the mail server. This means that you must put the full name of your mail server in this field. For instance use smtp.example.com (and not just "smtp"). The name must also be resolvable to an IP address using DNS. This means that the command "nslookup smtp.example.com" must work and must display the valid IP address of the machine on which Proto Balance Mail is running. Reverse DNS should also resolve back to the FQDN. Contact the administrators of your DNS server to get more information. For initial evaluation of Proto Balance Mail, it is not immediately necessary to change your DNS configuration.

Listen interface address: Ethernet IP on which to accept new TCP connections for email delivery. This will usually be 0.0.0.0 indicating to listen on *any* available interfaces. Occasionally however, if you have more than one network interface, you may choose to put the specific interface IP address here - this would be for security/firewalling purposes.

TCP listen port: The port on which to accept new connections. This is almost always port 25.

Comment: You can put arbitrary informational text here.

How to choose relay box: When a user of your email system attempts a remote delivery, there are two ways Proto Balance Mail can choose which box through which to relay that delivery. This can either be a random choice (the setting "Load balance over SMTP-relay-enabled boxes") or a choice based on the mail server configured for this user account (the setting "Use box in per-email config"). The mail server of the user account's "Box name" under the configuration for the particular senders email address. If you choose "Load balance over..." then you must ensure that at least one of your boxes has configured "SMTP relay: Enabled". If you use choose "Use box in..." then you must ensure that every email address in the email database has a "Box name" configured. The latter option is useful to ensure that the same box is used for both sending and receiving email w.r.t. a particular user.

Max mail size in kilobytes: This will cause an email to be rejected if it is over the size limit that you desire for your mail system. Note that this size limit *must* be *less* than that of any of your mail server boxes. For example, if you have three mail servers and their respective size limits are 10M, 20M, and 100M, then choose a value of 9000 (kilobytes) to be sure.

Allow notifications from relay clients: A common source of spam are desktop machines that send emails with an empty sender address. The setting "Disallowed" blocks this spam. This will also block notifications about email being read (some mail clients support such notifications) - however this is rarely a problem.

Default spam filtering: This is the setting that will be substituted for any user account that has a spam setting of "Default". Proto Balance Mail currently only supports "Greylisting" or "None". Note that Proto Balance Mail's greylisting a more advanced form generic greylisting which includes statistical analysis.

Sleep on certain SMTP errors: If a remote SMTP delivery behaves in a suspicious or erroneous fashion, Proto Balance Mail will artificially hold the connection open to appear to be "busy" for this number of seconds in order to frustrate the offending peer. You should set this value to reasonable random value of your choice.

Client relay address ranges: If your client desktop machines do not have passwords configured for their outgoing SMTP delivery, then you can authenticate them using their IP addresses. If you are an ISP, these will be the DSL address ranges that you dynamically serve. If you are a company or college, then these will be the local IP address ranges used on your campus.

Foreign SMTP server SPAM whitelist: These are address ranges of remote SMTP servers that should not ever be automatically blacklisted by Proto Balance Mail's auto-blacklisting mechanism. For safety you should include here the address ranges of the SMTP servers of the top providers of email services (like google, hotmail, yahoo, etc.)

Recipients per MAIL FROM: Desktop mail-bots will typically try to send a large number of emails using different sender addresses. Regular users normally use only a small number of personal email addresses. Proto Balance Mail counts the number of unique sender addresses used *per* IP for each of the IP addresses in the range of "Client relay address ranges". If this count exceeds the maximum you have specified, then the IP address is *blocked* and appears in the blocked list under "Show blocked IP addresses" in the "Email Database" tab. Note that this count applies only to sender addresses that are not in the email addresses database (i.e. "stranger" email addresses). Clients that exceed these limits are blocked with the error message "550 you have exceeded your mail quota - contact our site administrators for more info".

Spam spread factor: This feature is used to adjust Proto Balance Mail's proprietary suspect email analysis engine. Contact us before changing this setting.

Client block address ranges: This blocks SMTP deliveries outright from the specified address ranges. Proto Balance Mail will reply with the message "550 your service has been explicitly blacklisted - contact our site administrators for more info".

Box connect timeout in milliseconds: Default for new boxes. This setting does not affect already created boxes. See the same setting in the box configuration.

Maximum number of connections: Proto Balance Mail will drop new connections if the number of concurrent connections exceeds this setting. If you anticipate larger load than 2000 concurrent connections, see the generic user manual, "System Settings for High Performance", and use Proto Balance Mail with the -shared option.

Maximum number of per-client connections: Denial of service attacks might attempt to connect more often than is reasonable for a mail system. This setting can provide DoS protection.

Maximum number of per-client concurrent: Denial of service attacks might attempt more concurrent connections than is reasonable for a mail system. This setting can provide DoS protection.

15.2 General configuration settings

The settings under this section are under the "Info" tab.

The following settings relate to denial of service protection:

Delete an IP address from the dictionary: The IP address dictionary is a table of remote IP addresses; addresses from which Proto Balance received any kind of connection attempt. The table is used for analyzing connection rates to prevent denial of service (DoS) attacks. Specifically the "Maximum number..." settings in the cluster configuration relate to the statistics gathered in this table. If no further connection attempts come from an IP after the number of seconds indicated by this setting, then the entry is removed from the table.

Max dictionary size in megabytes: Proto Balance Mail saves memory by removing the oldest entries to keep the consumed memory below the amount specified here. See "Delete an IP address from the dictionary" above.

The following settings relate to spam-blocking imposed on foreign clients, meaning remote SMTP servers that attempt to deliver email to your local users:

Greylisting initial time: The period after the first delivery attempt for which the mail delivery is rejected with a "450" (temporarily unavailable) error. See the section "Spam Prevention".

Greylisting initial time for suspicious email: A longer initial denial period applied to email that appears to have the characteristics of spam.

Greylisting first window: The period wherein a second or subsequent delivery attempt will be allowed. If no delivery attempt happens by the end of this period, the greylisting record is deleted. If the delivery is allowed, the greylisting record becomes a permanent record.

Greylisting max record age: The storage period for permanent records. Some implementations set this to 36 days but because of Proto Balance Mail's compact storage of greylisting records, Proto Balance Mail can efficiently store a very large number of records. A value of 172800 (4 months) is not unreasonable. If you host a very large site and would like to reduce your memory consumption, then this value can be set to 51840 (36 days) or 11520 (8 days).

Whitelist max record age: Outgoing SMTP deliveries by authenticated clients are recorded in a whitelist. Proto Balance Mail consults the whitelist before deciding to greylist an incoming delivery. This means that if a desktop machine sends a message to an outside mail user, and that user replies, the reply will not be greylisted or otherwise delayed. This setting dictates how long to keep such records.

Spreadlist max record age: This is an internal setting for use with Proto Balance Mail's proprietary spam analysis engine. Please consult us before adjusting this setting.

The following settings relate to limits imposed on relay clients, meaning desktop machines or other systems that are allowed to use your Proto Balance Mail installation's outgoing delivery service. Such are identified either by being within "Client relay address ranges" or because they are authenticated using SMTP AUTH:

Recipient exceeded quota: The number of times a client can receive the error "number of recipients exceeded" before being blocked and listed under "Show blocked IP addresses".

Interval: The interval of the previous setting after which the internal count is reset. Spam bots tend to send a lot of email in a short period. Use an interval of between 1 and 10 minutes.

MAIL FROM total quota: The number of MAIL FROM commands a client can send before being blocked and listed under "Show blocked IP addresses".

Interval: The interval of the previous setting after which the internal count is reset. Spam bots tend to send a lot of email in a short period. Use an interval of between 1 and 10 minutes.

Distinct FROM emails per IP address: The number of unique sender addresses the client can use from the client's IP address. If you have clients that deliver mail from behind a masquerading firewall, and have a large number of desktop machines, this setting may need to be increased.

Interval: The interval of the previous setting after which the internal count is reset. Spam bots tend to use a lot of different sender addresses. Use an interval of between 60 and 600 minutes.

Distinct FROM domains per IP address: The number of unique sender domains that the client can use for the client's IP address. Similar to the previous setting but applying only to part of the address after the @ symbol.

Interval: The interval of the previous setting after which the internal count is reset. Spam bots tend to use a lot of different sender addresses. Use an interval of between 60 and 600 minutes.

15.3 Email and alias settings

The following settings relate to configuring user accounts. See under the "Email Database" tab. These settings pertain to one individual user's personal configuration:

Email address: This is the POP authentication login name, the SMTP authentication login name, and the email address of the local user account.

Password: This is password for POP authentication. It is also the SMTP AUTH authentication password for outgoing email. You need not configure a password if your users are not using POP or are POPping email directly from your back-end mail servers. For your users to send email, you need only add their IP address ranges to "Client relay address ranges" in the cluster configuration.

Box name: This is the back-end mail server on which the mailbox is stored. It must exactly match the name you have chosen for the box under the "Clusters" tab. The field is case-sensitive. If you are using Proto Balance Mail for outgoing email and you have configured "How to choose relay box" to the value "Load balance over...", then you can leave the "Box name" setting blank.

Return path equals recipient address: Because a lot of spam uses identical recipient and sender ("return path") addresses, Proto Balance Mail can, at the user's preference, explicitly reject such email by configuring this setting to "Disabled". This applies to incoming email only.

Null return path: Because a lot of spam uses null sender ("return path") addresses, Proto Balance Mail can, at the user's preference, explicitly reject such email by configuring this setting to "Disabled".

Spam filtering: Set the spam filtering for this user to one of "Disabled", "Default", or "Greylisting". The setting "Default" means to use the spam setting in the cluster configuration.

Recipients per MAIL FROM: This setting is the maximum number of unique recipients that a user is allowed to deliver email to in one delivery session. If the user exceeds this limit, the user's "recipient-exceeded" count is incremented. If this count exceeds "Recipient exceeded quota" (under the "Info" tab), then the user is blocked from further delivering email and appears in the blocked list under "Show blocked IP addresses". This will have the effect that spam bots running on your user's infected machines will not be able to send large amounts of email, thus preventing your mail servers from getting blacklisted.

Auto-reply subject: If a user wishes to enabled "out-of-office" auto-replies, then this setting can be configured. This setting accepts a Unicode UTF-8 string. Leave blank to disable auto-replies.

Auto-reply message body: See previous setting. This setting accepts a Unicode UTF-8 string. Leave blank to disable auto-replies. Auto-replies are send with "Content-Type" set to "text/plain; charset=UTF-8". This means that auto-replies are fully internationalized.

Auto-reply keep copy: If set to "Enabled", this setting causes email to continue to be delivered to the user's mailbox while also sending an auto-reply.

Auto-reply expiry date: The date and time up until which the auto-reply will be active.

Unlimited email: If set to "Enabled", this will allow the user to send an unlimited amount of email. See next setting.

Unlimited email expiry date: This will set a date limit after which the unlimited email setting will be disabled. The settings "Unlimited email" and "Unlimited email expiry date" are useful for allowing a user to run a mailing list which will send out large amounts of legitimate email for a specific time.

Remote deliveries: A significant amount of spam uses your own user addresses as sender addresses. By configuring this setting to "Disallowed", any email delivered from a remote location, that maliciously uses this user's email address as a sender address, will be blocked. This has the disadvantage that this user will not be able to send email from a foreign SMTP service to someone at the same site. For example the user will not be able to send email from the user's BlackBerry if one of the recipients is also customer of yours. Therefore use with caution.

16 Spam Prevention

Proto Balance Mail implements the SPAM blocking technique of "grey-listing" (or Gray-listing or Greylisting). Grey-listing is today supported by most mail servers. We recommend reading the Wikipedia entry:

`http://wikipedia.org/Grey_listing`

Grey-listing requires that your organization have a single database of grey-list records. This means that a traditional cluster of mail servers cannot effectively implement grey-listing if the nodes in the cluster separately implement the grey-list algorithm. Proto Balance Mail solves this problem by implementing grey-listing in one place.

The use of grey-listing does not prohibit the use of more processor intensive SPAM analysis software on your back-end mail server. Indeed you may continue to run SPAM and virus filters on all your mail servers.

Proto Balance Mail varies the grey-listing algorithm slightly. Grey-listing defines an initial period during which an email is blocked. This period defaults to "Greylisting initial time" (under the [Info] tab), which is usually set to 5 minutes. Should the email appear suspicious, however, this initial period may be automatically increased to as much as "Greylisting initial time for suspicious email" (under the [Info] tab), which is usually set to 2 hours.

(The precise algorithm by which email is judged to to suspicious cannot be disclosed for obvious reasons. IP addresses over a certain threshold of "suspicious" are blacklisted outright.)

If the remote mail server performs a delivery before after the initial period and before "Greylisting first window" (usually 6 hours) the email will be permitted and become a "permanent record". After 6 hours the record will expire and the algorithm will start again.

Permanent records are kept for "Greylisting max record age" - usually 120 days.

17 Spam Blocking Effectiveness

The preceding techniques resulted in SPAM being reduced by a factor of 50 on one of our customers sites. However, the effectiveness of the technique depends on the scale of your email cluster. A 98% reduction was achieved on a system of 20,000 active mail boxes. However the blocking effectiveness reduces with a smaller number of mail boxes.

In general, the *larger* your operation then the *better* is Proto Balance at blocking SPAM.

18 XML Interface (SOA)

Proto Balance Mail can be accessed programmatically using a Service Oriented Architecture (SOA) interface based on XML. Actions such as adding a new mail box, configuring alias address settings, changing passwords, and adjusting box/cluster settings: these actions can all be done remotely using a programming language of your choice.

Please contact us for a complete repertoire of example code in Java, Perl, Python, and other programming languages.

All XML actions are executed using a POST request to the URL:

```
http://1.23.45.67:8080/xml
```

All XML content is of the form:

```
<?xml version="1.0" encoding="UTF-8" ?>
<xxx>
        <accesskey>qwerty</accesskey>
        <yyy>zzz</yyy>
        ...
</xxx>
```

Where the "accesskey" tag is that entered under Info - Change Password. A tracedump of a simple XML post, say to change a user's password, would be:

```
POST /xml HTTP/1.0
content-length: 224

<?xml version="1.0" encoding="UTF-8" ?>
<commitmodifyemail>
        <accesskey>qwerty</accesskey>
        <email_address>abigail@example.com</email_address>
        <user_password>45b4a9216f9f</user_password>
</commitmodifyemail>
```

19 XML Command Reference

The following reference includes XML actions to manipulate the user account database. Further XML examples - for example to modify the cluster and box configuration - are available on request.

Change a user's password:

```
<?xml version="1.0" encoding="UTF-8" ?>
<commitmodifyemail>
        <accesskey>qwerty</accesskey>
```

```
        <email_address>abigail@example.com</email_address>
        <user_password>Qrlhax7IMF</user_password>
</commitmodifyemail>
```

Change a user's backend mail-server:

```
<?xml version="1.0" encoding="UTF-8" ?>
<commitmodifyemail>
        <accesskey>qwerty</accesskey>
        <email_address>abigail@example.com</email_address>
        <boxid>MAILSERVER002</boxid>
</commitmodifyemail>
```

Change a user's complete configuration:

```
<?xml version="1.0" encoding="UTF-8" ?>
<commitmodifyemail>
        <accesskey>qwerty</accesskey>
        <email_address>abigail@example.com</email_address>
        <user_password>Qrlhax7IMF</user_password>
        <boxid>MAILSERVER001</boxid>
        <to_equals_from>1</to_equals_from>
        <null_reverse_path>1</null_reverse_path>
        <spam_setting>1</spam_setting>
        <rcpt_per_mail>50</rcpt_per_mail>
        <auto_reply_subject></auto_reply_subject>
        <auto_reply_body></auto_reply_body>
        <auto_reply_copy>1</auto_reply_copy>
        <auto_reply_expiry></auto_reply_expiry>
        <mail_unlimited></mail_unlimited>
        <mail_unlimited_expiry></mail_unlimited_expiry>
        <allow_remote_deliveries>1</allow_remote_deliveries>
</commitmodifyemail>
```

Retrieve a user's complete configuration:

```
<?xml version="1.0" encoding="UTF-8" ?>
<getemail>
        <accesskey>qwerty</accesskey>
        <email_address>abigail@example.com</email_address>
</getemail>
```

Change an alias's complete configuration:

```
<?xml version="1.0" encoding="UTF-8" ?>
<commitmodifyalias>
        <accesskey>qwerty</accesskey>
        <alias_address>sales@example.com</alias_address>
        <alias_list>abigail@example.com,yolanda@example.com</alias_list>
        <rcpt_per_mail>50</rcpt_per_mail>
        <mail_unlimited></mail_unlimited>
        <mail_unlimited_expiry></mail_unlimited_expiry>
        <spam_setting>1</spam_setting>
        <allow_remote_deliveries>1</allow_remote_deliveries>
</commitmodifyalias>
```

Retrieve an alias's complete configuration:

```
<?xml version="1.0" encoding="UTF-8" ?>
<getalias>
        <accesskey>qwerty</accesskey>
        <alias_address>sales@example.com</alias_address>
</getalias>
```

Unblock a user and/or IP address:

```
<?xml version="1.0" encoding="UTF-8" ?>
<setblockip>
        <accesskey>qwerty</accesskey>
        <block_email>0</block_email>
        <email_address>abigail@example.com</email_address>
        <address>1.2.3.4</address>
</setblockip>
```

Retrieve a list of blocked IPs: This is a list of IP addresses of your own users who have been blocked due to mailware/viruses sending mail through Proto Balance):

```
<?xml version="1.0" encoding="UTF-8" ?>
<blockedips>
        <accesskey>qwerty</accesskey>
</blockedips>
```

Delete an email address:

```
<?xml version="1.0" encoding="UTF-8" ?>
<deleteemail>
        <accesskey>qwerty</accesskey>
        <email_address>abigail@example.com</email_address>
</deleteemail>
```

Delete an alias address:

```
<?xml version="1.0" encoding="UTF-8" ?>
<deletealias>
        <accesskey>qwerty</accesskey>
        <email_address>sales@example.com</email_address>
</deletealias>
```

Add a new email address:

```
<?xml version="1.0" encoding="UTF-8" ?>
<addnewemail>
        <accesskey>qwerty</accesskey>
        <email_address>abigail@example.com</email_address>
        <user_password>Qrlhax7IMF</user_password>
        <boxid>MAILSERVER001</boxid>
        <to_equals_from>1</to_equals_from>
        <null_reverse_path>1</null_reverse_path>
        <spam_setting>1</spam_setting>
        <rcpt_per_mail>50</rcpt_per_mail>
        <auto_reply_subject></auto_reply_subject>
        <auto_reply_body></auto_reply_body>
        <auto_reply_copy>1</auto_reply_copy>
        <auto_reply_expiry></auto_reply_expiry>
        <mail_unlimited></mail_unlimited>
        <mail_unlimited_expiry></mail_unlimited_expiry>
```

```
        <allow_remote_deliveries>1</allow_remote_deliveries>
</addnewemail>
```

Add a new alias address:

```
<?xml version="1.0" encoding="UTF-8" ?>
<addnewalias>
        <accesskey>qwerty</accesskey>
        <alias_address>sales@example.com</alias_address>
        <alias_list>abigail@example.com,yolanda@example.com</alias_list>
        <rcpt_per_mail>50</rcpt_per_mail>
        <mail_unlimited></mail_unlimited>
        <mail_unlimited_expiry></mail_unlimited_expiry>
        <spam_setting>1</spam_setting>
        <allow_remote_deliveries>1</allow_remote_deliveries>
</addnewalias>
```

20 Frequently Asked Questions

- Can I deploy multiple Proto Balance Mail machines in a "no-single-point-of-failure" arrangement?

  Yes. Multiple Proto Balance Mail machines will discover each other and keep their user account databases and client access information replicated and synchronized. See "Multiple Proto Balance Mail Machines" in the Mail User Manual.

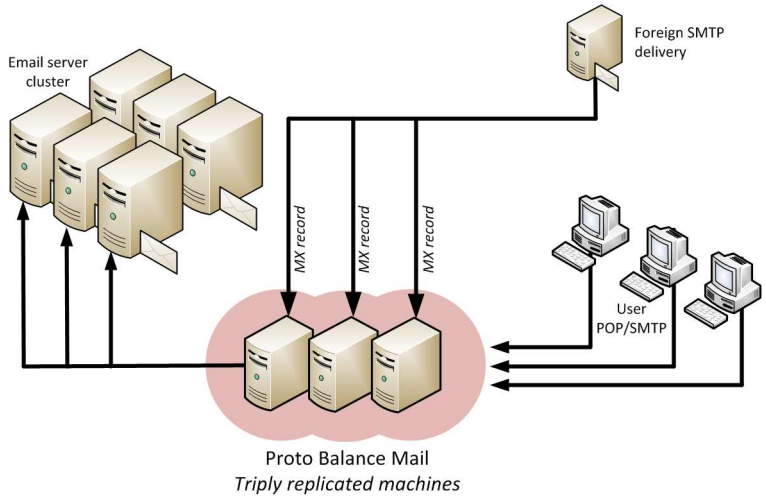

- How do Proto Balance Mail instances discover each other?

  Proto Balance Mail broadcasts UDP packets on port 7931 and 7932 using a discovery and replication protocol. Proto Balance Mail machines will automatically negotiate sessions with all other machines on the local LAN as well as all machines explicitly configured as sister hosts. Protocol packets are crypto-signed for security.

- How many Proto Balance Mail machines can I have replicating between one another?

  Proto Balance Mail replicates all data to all discovered sister machines. The replication protocol is a binary encoded protocol of very small packets and is therefore light-weight and fast. However network load will increase quadratically with the number of Proto Balance Machines. Therefore, we recommended proceeding cautiously over 20 machines. Note that this has nothing to do with the number of back-end mail servers - which is *unlimited*.

- If I have a University/College with different faculties, can I put each faculty on a different mail server even if all email addresses have the same domain part of the email address?

  Yes.

  You can have separate email servers for each faculty or department, that do not have to know about each other nor share the same file-system. Only Proto Balance Mail will know about all of them.

- Does Proto Balance Mail work with Microsoft Exchange, qmail, PostFix, sendmail?

  Yes. Proto Balance Mail strictly implements the SMTP and POP protocols. All SMTP servers are supported.

- Is Proto Balance Mail an MTA? What is an MTA?

  MTA means Mail Transfer Agent. This is an email server (like Microsoft Exchange or sendmail) that receives email, queues it on disk, and appends it onto local mailboxes of users who have accounts within.

  Proto Balance Mail is not this. It does not store email but directs email traffic (SMTP) to the server suited to receive delivery. Proto Balance Mail can do this even if there are ten or a hundred MTAs connected directly to it.

- How do I *gradually* migrate to Proto Balance Mail?

  You begin by installing it on a separate machine and following the instructions in "Mail User Manual". However you do not need to migrate your entire system. You can begin by testing a single user.

  With some mail solutions you have to migrate all your mail servers: an all-or-nothing affair that is risky. With Proto Balance Mail, you perform an incremental migration where each step is risk-free and reversible.

  Consider if you have a typical mail server deployment as follows:

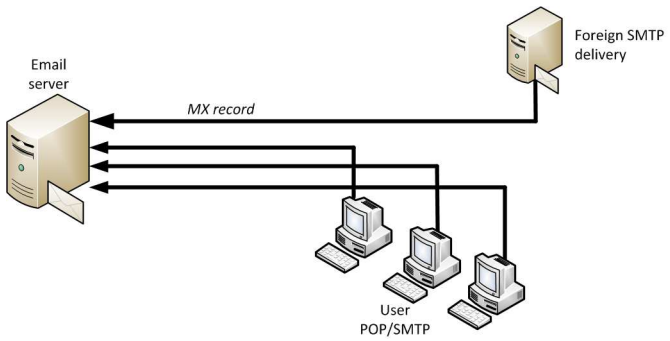

  Begin by deploying Proto Balance Mail and moving a single trial user:

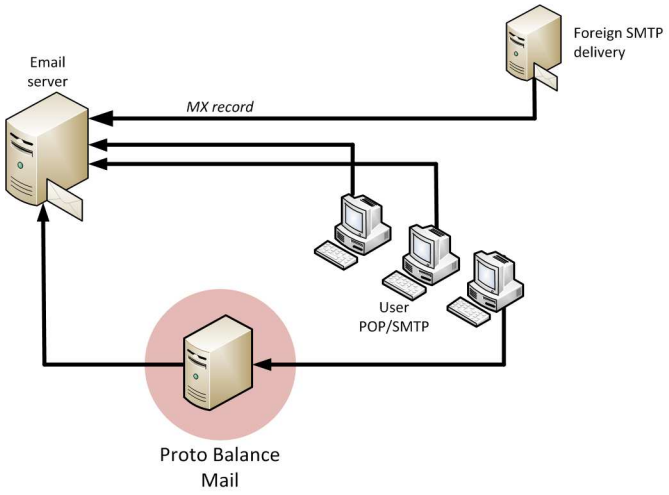

Next gradually migrate all your users as follows:

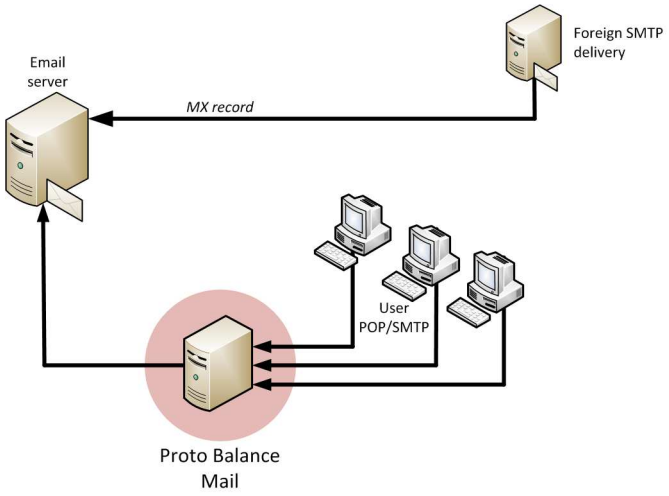

Finally, migrate your MX record of your DNS (Domain Name Server):

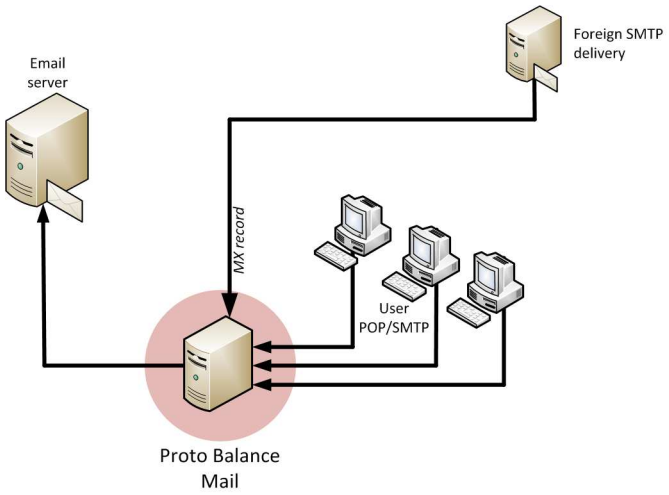

Your system will now work as before, but with Proto Balance Mail's advanced Spam filtering system.

Now you can scale your new system, apportion your mailboxes, divide users into groups as you please:

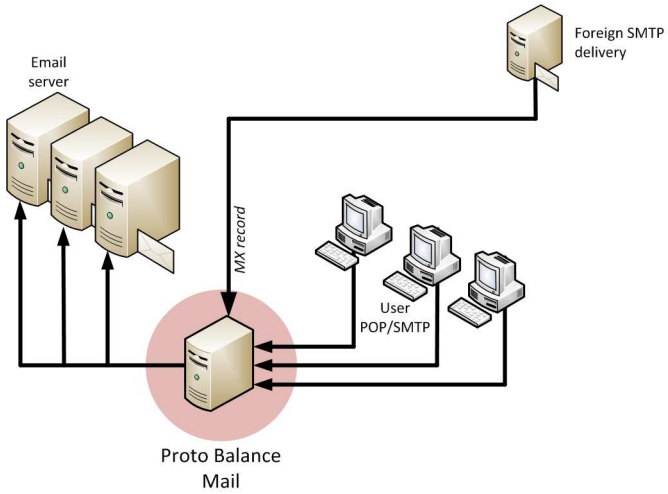

- Why is this better than any other system?

  Because no other clustering solution allows you to apportion mailboxes to separate machines. With Proto Balance Mail you can, for example, put your sales department on one mail server, and your engineering department on another mail server, *even* if those departments have the identical domain part of the email address (e.g. @mycorp.com). Without Proto Balance Mail this is only possible by changing the domain (e.g. @sales.mycorp.com, @eng.mycorp.com).

  If you are an ISP you can put different geographic regions on different mail servers. In fact you can apportion the mailboxes any way you like. And relocate them whenever you like. Typically, Proto Balance Mail customers put 8000 mailboxes to a mail server.

- How does Proto Balance Mail achieve such high traffic throughput?

  Most SMTP forwarding software receives the entire of the email, storing it in a queue directory on your hard disk. Proto Balance Mail does not store the email, but directs the SMTP session on-the-fly to your email server.

  It does this using a patented method that is also able to create concurrent duplicate sessions. This means alias expansion works.

- How do I setup Proto Balance Mail just for blocking of spam?

  If you are not interested in the load balancing features and just want to use Proto Balance Mail's sophisticated spam blocking functionality; then you can channel just incoming email through Proto Balance Mail.

  Consider if you have a three-mailserver system as follows:

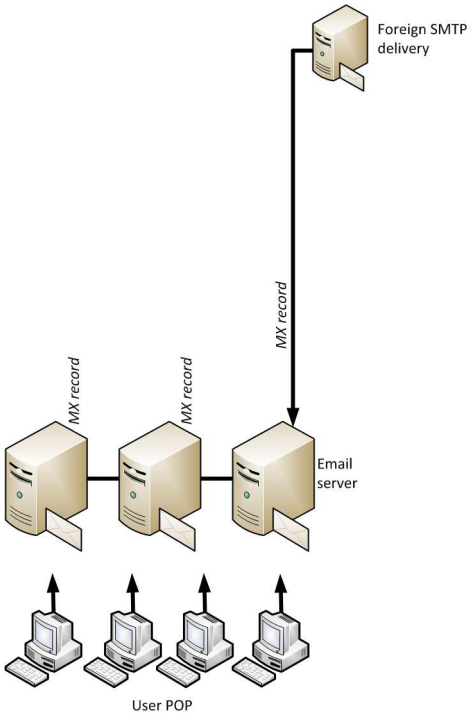

Begin by installing Proto Balance Mail. For Proto Balance Mail's spam-blocking to work, you need to create a flat file of all email addresses, as well as all aliases. Follow the instructions in the "Mail User Manual". You will need to randomly assign a box to each email record so that one third of the deliveries direct to each mail server. Now deploy Proto Balance Mail and create a cluster with boxes pointing to each of your mail servers. Setup each box to be a fall-over to its neighboring box:

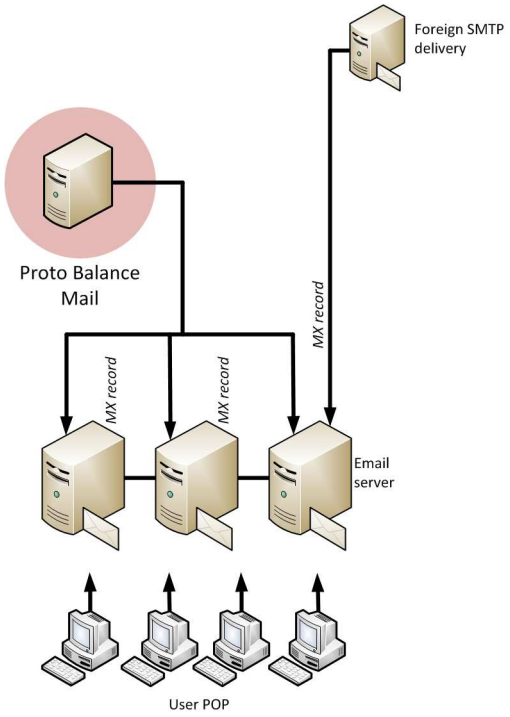

You can now test that mail delivery works, either by setting up a dummy delivery from a foreign SMTP server (tricked to thinking your MX record is that of Proto Balance Mail), or you can telnet

to Proto Balance Mail and issue a mail session manually (see the "Mail User Manual" for how to do this). Note that grey-listing will block first deliveries - you need to wait until the grey-listing first interval has expired and try again. Alternatively, you can turn off grey-listing by setting "Default spam filtering" to "Disabled" in the cluster configuration.

Finally you can switch over your MX records:

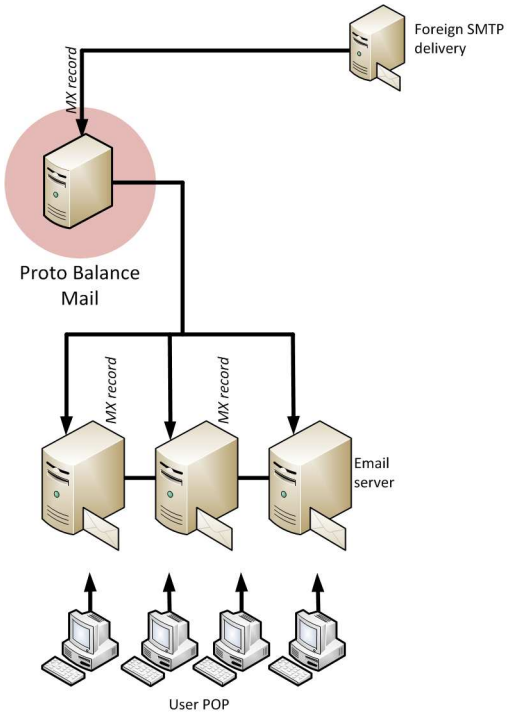

Note that it is not ideal to use Proto Balance Mail for incoming spam filtering without also doing outgoing deliveries through Proto Balance Mail. This is because outgoing deliveries create white-list records that prevent incoming email from being delayed. However it is still an acceptable solution.

- How do I load balance an outgoing mail queue with Proto Balance Mail?

  This question pertains to admins who want to use Proto Balance Mail for outgoing mail and are not interested in POP or incoming SMTP. You would be asking this question if you found your outgoing SMTP servers experiencing high load and very long queue lengths. There are two ways to load balance outgoing SMTP:

  The preferred way is to use Proto Balance Mail's capabilities to direct SMTP sessions. Begin by adding an "SMTP Cluster" in the config web page. Add "Boxes" for each outgoing mail server and set their traffic lights to green. In the cluster configuration set "Client relay address ranges" to cover all the IP addresses that are allowed to relay email. Then set "Recipients per MAIL FROM..." to a high value like 10000. Set "Allow notifications..." at your preference. Set "How to choose relay box" to [Load balance over SMTP-relay-enabled boxes]. Under the "Info" tab, set the four settings "Distinct FROM emails/domains per IP address", "MAIL FROM quota", and "Recipient exceeded quota" to a high values like 10000. Then go to each box and select "SMTP relay" to [Enabled].

  Such settings will result in outgoing SMTP deliveries being more-or-less randomly distributed between the boxes. Proto Balance Mail will use "consistent hashing" based on the sender address to choose the box. This means that the same sender will tend to send through the same box unless that box is unavailable.

  The second way to load balance is to use Proto Balance Mail's pure TCP load balancer with queue monitoring. Click on "Add Cluster" to add a pure TCP load balancer listener on port 25. Add a Box for each outoging SMTP server. Set all traffic lights to yellow. You now need to write scripts

that run on your SMTP server that can turn the box traffic lights from green to yellow based on the mail queue length:

Begin by going to the "Login" tab and clicking on [Change password]. Here you will be able to enter your XML access key. Let us say that your XML access key is "qwerty12345" and your cluster is called "Cluster001". In order to change BOX001's traffic light to green, your SMTP server needs to issue the following XML to port 8080 of Proto Balance Mail:

```
<?xml version="1.0" encoding="UTF-8" ?>
<green>
        <accesskey>qwerty12345</accesskey>
        <clusterid>Cluster001</clusterid>
        <boxid>BOX001</boxid>
</green>
```

If your mail queue becomes too full, you can issue the XML:

```
<?xml version="1.0" encoding="UTF-8" ?>
<yellow>
        <accesskey>qwerty12345</accesskey>
        <clusterid>Cluster001</clusterid>
        <boxid>BOX001</boxid>
</yellow>
```

An example python script to do this is:

```
import socket, string
from myconfig import *

balance_address = "192.168.0.1"

def xmlsubmit(xml):
        c = socket.socket()
        c.connect ((balance_address, 8080))
        c.send("POST /xml HTTP/1.0\r\n")
        c.send("content-length: %s\r\n" % len(xml))
        c.send("\r\n%s" % xml)
        r = ""
        while 1:
                v = c.recv(4096)
                if not v:
                        break
                r = r + v
        return r

red_yellow = """
        <?xml version="1.0" encoding="UTF-8" ?>
        <yellow>
                <accesskey>qwerty12345</accesskey>
                <clusterid>Cluster001</clusterid>
                <boxid>BOX001</boxid>
        </yellow>
"""

print xmlsubmit(red_yellow)
```

How you establish the length of your mail queue depends on the SMTP software you have available.

- Will Proto Balance Mail handle email alias addresses?

  Yes.

  With Proto Balance Mail you can configure aliases to different users, even foreign email addresses and across different mail servers. Proto Balance Mail will transparently copy all emails in the alias expansion.