



Managing mail traffic with an IMAP proxy

# MESSAGE HANDLER

IMAP proxies like Perdition, Imapproxy, and Cyrus Aggregator help distribute mail to multiple IMAP servers. We examine some options for IMAP proxy in the Linux environment. **BY JACK CHONGJIE XUE AND MARKUS FEILNER**

**M**ore and more IMAP servers are groaning under the load of increasing mail traffic. This presents administrators with a problem: clustering an IMAP server is not a trivial task. To make sure a shared database is available for user mailboxes, administrators often resort to expensive SANs for shared storage or complex replication mechanisms.

The cost of commercial clustering systems such as Red Hat's cluster suite RHCS make some IT bosses nervous, and the prospect of migrating an old IMAP server to a new cluster system and the technical complexities involved bring many a system administrator to a nervous sweat.

Linux admins who set up their own systems based on free software are typi-

cally frustrated by the mediocre high-availability capabilities of the current crop of IMAP servers. While redundant solutions for the MTAs (Mail Transfer Agents) of today's popular SMTP servers are easily achieved by DNS or IP-based load balancing techniques, the quality of the storage back end determines the success or failure of an IMAP cluster.

## IMAP: Easier than Clustering

An IMAP proxy is a simple and flexible alternative to a clustering solution. IMAP proxies separate message transport and delivery from data storage. Email clients do not use IMAP or POP to talk to the mail server but simply communicate with the proxy. The IMAP proxy then checks its database, finds the IMAP

server responsible for the current user, and passes the request on to the responsible server.

## Front Ends for IMAP

You can think of an IMAP proxy as a front end for the mail server. The client connects to the proxy, and the proxy communicates with back-end mail servers to obtain the user's mail data. This approach adds unexpected flexibility, as the client always connects to the same host, even if the mail data has moved to a completely different location.

Standard IMAP protocol features offer many forwarding and redirection functions; however, the various software projects use different approaches to perform what is basically the same task. All the IMAP proxy tools have two basic things in common:

- The front end and back end must match and provide the same IMAP functionality.
- The proxy uses a database to manage user-server mappings; access is local or via the network.

One of the best-known and most flexible IMAP proxies is Perdition [1]. Another IMAP proxy called Imapproxy [2] focuses on the problem of webmail clients. Server-specific packages such as Cyrus IMAP Aggregator [3] offer a huge collection of features, but typically only for the matching IMAP server.

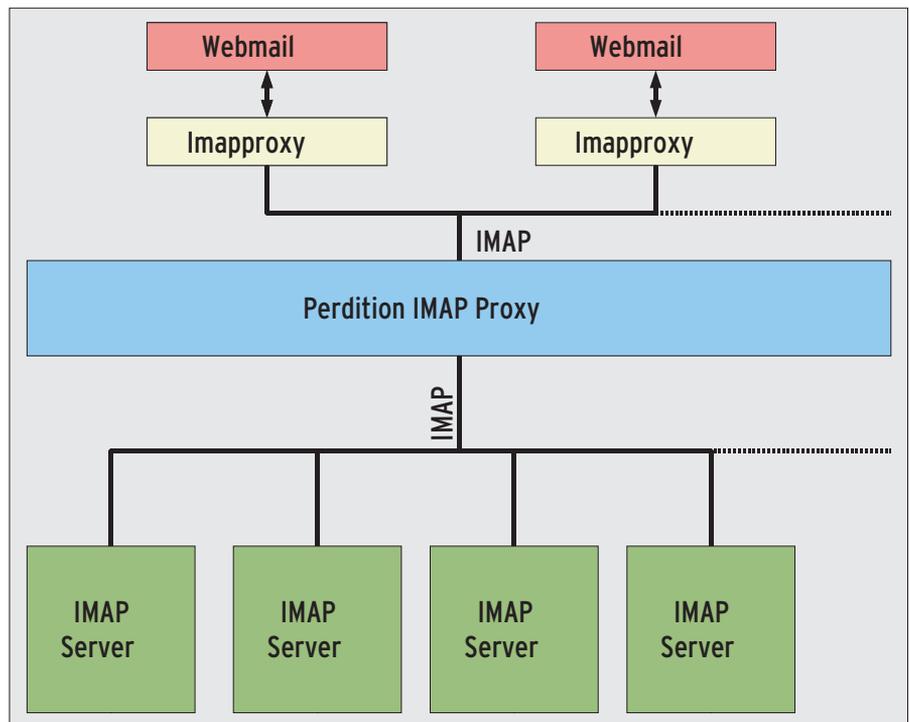
## Perdition

Perdition is a simple mail retrieval proxy that supports the IMAP4 and POP3 standards through optional SSL or TLS connections. Perdition only understands the IMAP commands it needs for authentication purposes; otherwise, it just passes data through as is.

The “Perdition Paper” by Simon Horman [4] has more details on the features and workings of Perdition for those readers who are interested in a more in-depth view. Even though this document has been around for a couple of years, it is still regarded as *the* reference for Perdition.

The IMAP proxy issues a database request and learns the correct server for the incoming client connection from the response.

One notable feature of Perdition is its support of multiple data sources. Perdition can speak to LDAP, ODBC, MySQL, and PostgreSQL back ends, or a mixture of all of these sources if necessary. Another advantage of the Perdition IMAP



**Figure 1:** While Perdition is a full-solution IMAP and POP proxy, Imapproxy mainly helps to accelerate webmailers by providing a local IMAP connection cache.

proxy is that it is modularly extensible.

## Installation

Perdition is very easy to install. Linux users have a choice of either downloading the source code and building from scratch, or using the binary RPMs. If you are interested in extended logging, you can additionally install the Vanessa Log-

ger [5]. Listing 1 shows the configuration for Perdition with a lookup table in a MySQL database. In this example, the database is located on the local machine; Listing 2 shows its contents.

During the migration, automated scripts can copy mailboxes from the old server to the new server and then update the necessary database entries. The Perdition proxy automatically redirects migrated users to the new back-end server. No changes are required on the client side, and users can enjoy uninterrupted email service.

## Imapproxy

Imapproxy is designed for the highly specialized task of accelerating webmail server performance by caching POP and IMAP server connection data.

By serving as an IMAP cache, Imapproxy reduces the number of connections required by a webmail client. Normally, each action by the webmail client triggers a read or write operation on the IMAP server, thus requiring the exchange of several data packets. Imapproxy's cache can handle many requests directly, thus accelerating the webmail application.

In contrast to Perdition, Imapproxy is more of an optimization tool than a genuine proxy. Imapproxy is useful wher-

### Listing 1: /etc/perdition/perdition.conf

```
01 connection_logging
02 connection_limit 1000
03 map_library /usr/lib/libperditiondb_mysql.so.0
04 map_library_opt "localhost:3306:dbPerdition:tblPerdition:perdition:
  x:username:servername:port"
05 username_from_database
```

### Listing 2: MySQL-Table

```
01 mysql> select * from dbPerdition.tblPerdition;
02 +-----+-----+-----+-----+
03 | user | servername | port |
04 +-----+-----+-----+-----+
05 | name1 | imap01.company.com | 143 |
06 +-----+-----+-----+-----+
07 | name2 | imap02.company.com | 143 |
08 +-----+-----+-----+-----+
09 ...
```

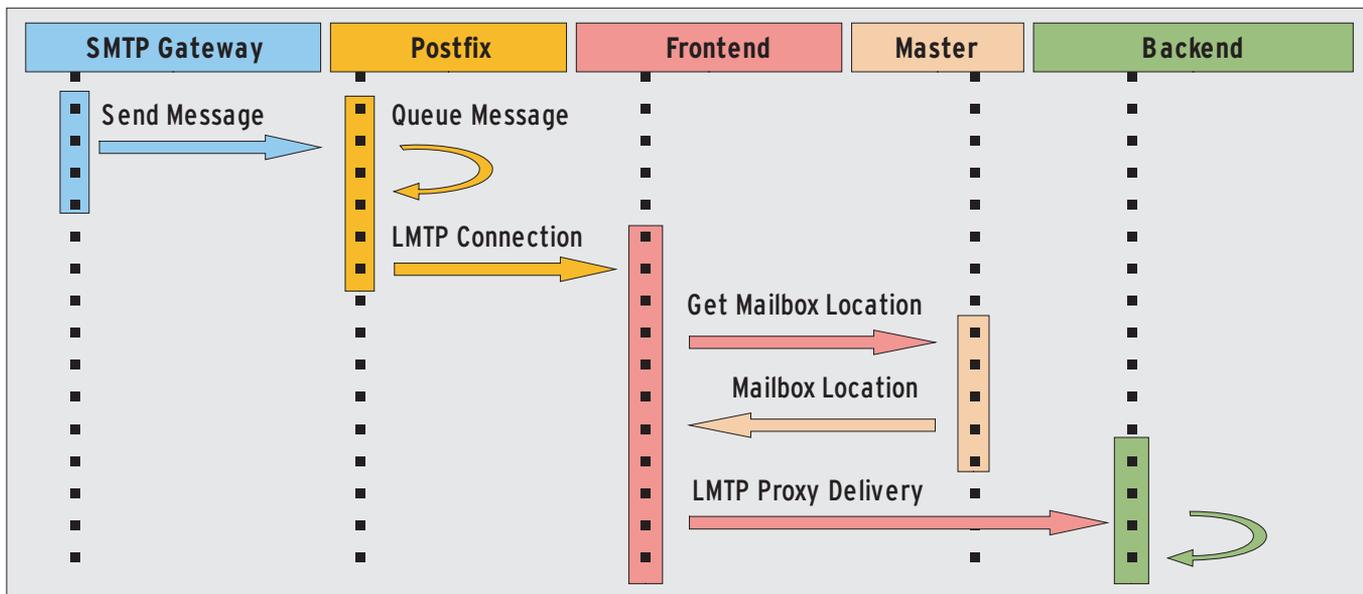


Figure 2: The mail gateway delivers the incoming message to the SMTP server. The SMTP server uses LMTP to forward the message to the front-end server, and finally the message is delivered to the matching back end.

ever users need browser-based access to the IMAP server. You can even combine the Imapproxy with a Perdition proxy (Figure 1).

### High Performance with the Cyrus Back End

Another IMAP proxy solution is Cyrus IMAP Aggregator Proxy (or Cyrus IMAP Proxy for short). The Cyrus proxy is probably the most robust tool of its kind. IMAP Aggregator Proxy supports both mail retrieval (POP/IMAP) and mail delivery (LMTP), although only for the Cyrus mail server.

Along with UW, Courier, and Dovecot, Cyrus IMAP is one of the big four IMAP servers, and it probably offers the biggest variety of features.

Cyrus stores email in its own format on the server, completely ignoring the tried-and-trusted mailbox or maildir formats, and it does its own thing in multi-

server setups. The Cyrus Murder Aggregator supports high-performance and load-balancing clusters by setting up front-end servers to communicate with storage back ends. Pairwise replication and shared storage via clustering filesystems like GFS provide ample opportunity for tinkering. But all of this takes time, tenacity, and money.

Aggregator is often used as an IMAP proxy for the popular Cyrus mail server. In this scenario, multiple back ends share the work, and the front end distributes the requests. Details of which user mail account resides on which server are stored on a separate machine known as the Mupdate server. The server acts as a kind of master node and uses the Mupdate protocol [6] to hand out data required for SMTP, IMAP, or POP requests.

### Aggregator as a Mail Delivery Proxy

Apart from mail delivery, Cyrus Aggregator is similar to Perdition; however, it uses its own server and a modified protocol. The big advantage of Aggregator is that it also supports SMTP and LMTP; thus, it provides a complete mail proxy system in collaboration with Cyrus (but only with Cyrus).

### Conclusions

With very little effort, an IMAP proxy can add functionality to an existing environment that would otherwise require a

more expensive clustering scenario. All of the programs described in this article make it easier for administrators to handle the daily flood of mail.

The IMAP proxy options for Linux add convenience and improve performance for busy mail systems. If you are planning to migrate your server and you are looking for compatibility, Perdition is the right choice.

Imapproxy fulfills its job as a cache for webmailers; it is not a competitor to Perdition as a full solution, but it is a useful extension. Cyrus Aggregator can distribute mail delivery over multiple back-end servers. ■

INFO	
[1] Perdition:	<a href="http://www.vergenet.net/linux/perdition">http://www.vergenet.net/linux/perdition</a>
[2] IMAP proxy:	<a href="http://www.imapproxy.org">http://www.imapproxy.org</a>
[3] Cyrus IMAP:	<a href="http://cyrusimap.web.cmu.edu/ag.html">http://cyrusimap.web.cmu.edu/ag.html</a>
[4] Perdition Paper:	<a href="http://www.vergenet.net/linux/perdition/perdition_paper/html">http://www.vergenet.net/linux/perdition/perdition_paper/html</a>
[5] Vanessa Logger:	<a href="http://www.vergenet.net/linux/vanessa/download/vanessa_logger">http://www.vergenet.net/linux/vanessa/download/vanessa_logger</a>
[6] "Scaling up Cambridge University's email service":	<a href="http://www.uxsup.csx.cam.ac.uk/~fanf2/hermes/doc/talks/2004-02-ukuug/paper.html">http://www.uxsup.csx.cam.ac.uk/~fanf2/hermes/doc/talks/2004-02-ukuug/paper.html</a>

THE AUTHOR

"Jack" Chongjie Xue works at Marshall University in Huntington, West Virginia, where he recently migrated the email cluster used by 30,000 students from Open VMS to Linux without any downtime. Besides Linux administration, Jack is interested in open source business intelligence, data visualization, and mathematics.

