

The sys admin's daily grind: Archivemail

BOX OF LEGENDS

Charly loves to be organized, but he also likes to have access to mail that reached him when the dinosaurs were still roaming the earth.

BY CHARLY KÜHNAST



```
File Edit View Terminal Help
charly@kintyre:~$ archivemail --days=90 -o /home/charly/mailarchive/ \
/home/charly/mailbox
/home/charly/mailbox:
  archived 39678 of 62728 message(s) (561.0MB of 865.3MB) in 154.9 seconds
charly@kintyre:~$ ls -lh *.gz
-rw----- 1 charly users 95M 2009-10-04 17:07 mailbox_archive.gz
```

Figure 1: In this example, Archivemail takes messages of 90 days or older from an overflowing mailbox and archives them in a Gzip file.

Occasionally I need to delve into my mailboxes for an antique message, just to prod my memory. Because my mail client has a full text index of the inbox, I usually find the information I'm looking for within a couple of seconds. Unfortunately, this flush of success is often a delicate thing: With 300 messages flooding in every day, not counting mailing lists, the indexer often gives up the fight. And because many messages have bulky attachments, I genuinely need to clean out my modern day equivalent of the Augean stables.

Archivemail [1] can help me clean up my mailbox; in fact, it has a couple of neat options for doing so. The *dry run* switch makes sure my mail survives the initial experiments. After setting it, Ar-

chivemail simply simulates the process without actually writing. The simple command line

```
archivemail --dry-run \
/home/charly/mailbox
```

packs messages that have lived in the mailbox for more than 180 days into a separate *mailbox_archive* file, then unbloats the file with Gzip.

Daily

Archivemail creates a new mbox-formatted file so that I can easily read the file with a mail client if I need to do so. Of course, you can configure the expiry date: If I want to get rid of mail before the 180-day limit, I could specify half this value,

```
archivemail --dry-run --days=90 \
/home/charly/mailbox
```

as shown in Figure 1. But in fact, I prefer to keep my messages far longer: *days = 730* to be precise. To keep my inbox lean despite this, I need to call Archivemail again:

```
archivemail --dry-run --size=1000000 \
/home/charly/mailbox
```

This sends all the messages greater than 1MB in size to the archive.

Archivemail has a couple of other useful options. If I use the archive mbox regularly or need permanent access, *--no-compress* prevents Gzip compression. If you get behind and can't keep up with all your messages, *--preserve-unread* excludes unread mail from the archiving process, even if time- or size-based criteria normally would apply.

The one thing you really need to organize your life is a big trash can. This adage also applies without restriction to mailboxes. If you are of the same mind, you might also want to try the *--delete* switch. This moves any messages that match the search criteria directly to the trash can rather than going through the archiving step. ■

INFO

[1] Archivemail: <http://archivemail.sourceforge.net>

SYSADMIN

Security Lessons 58

Rate limiting lets you deny service to attackers while legitimate users work on.

mod_selinux 60

Protect your web-based applications with Apache's mod_selinux module.

THE AUTHOR

Charly Kühnast is a Unix operating system administrator at the Data Center in Moers, Germany. His tasks include firewall and DMZ



security and availability. He divides his leisure time into hot, wet, and eastern sectors, where he enjoys cooking, fresh water aquariums, and learning Japanese, respectively.