

DOVECOT at DESY

First experiences with DOVECOT as IMAP server

Wolfgang Friebe
Dirk Jahnke-Zumbusch
HEPiX Spring 2009
Umeå University, Sweden

Reasons to change the IMAP server

- > at both DESY sites the hardware had to be renewed
→ time for (re-) thinking about requirements
- > problems / deficiencies with the old servers: mostly mbox related
 - mbox file format for INBOX (and all other mailboxes)
 - the number of big e-mails accumulate → mbox files are getting huge (see below)
 - DESY had mailboxes reaching the UW-IMAP and Solaris 8 UFS file limits (2GB)
 - only one process may write to mbox

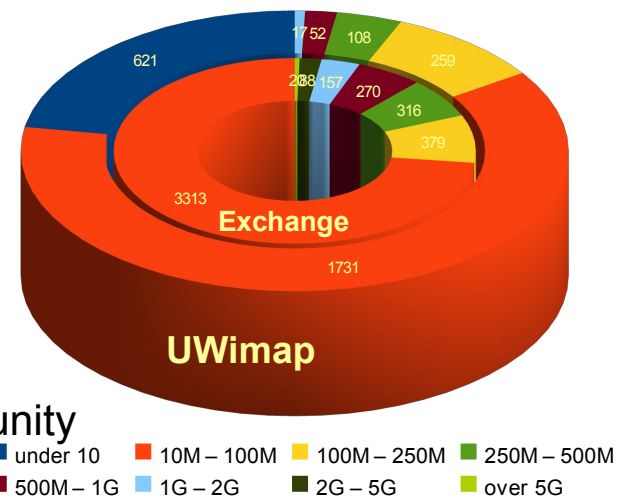
> requirements / needs

- server-side e-mail filtering
- access control for shared folders and mailboxes
- general speed up

> Dovecot looks promising

- Fast progressing development with very active community
- one of the fastest open source IMAP servers

mailbox sizes
distribution of mailboxes by size



Dovecot highlights

> Maildir++ INBOX and folder format

- fast access
- no locking issues

> plug-ins for enhanced functionality

- ACL support (local as well as via IMAP commands)
- compressed folder
- quota handling (local and via IMAP commands)
- ~20 plug-ins available (<http://wiki.dovecot.org/Plugins>)

> Namespaces easily configurable

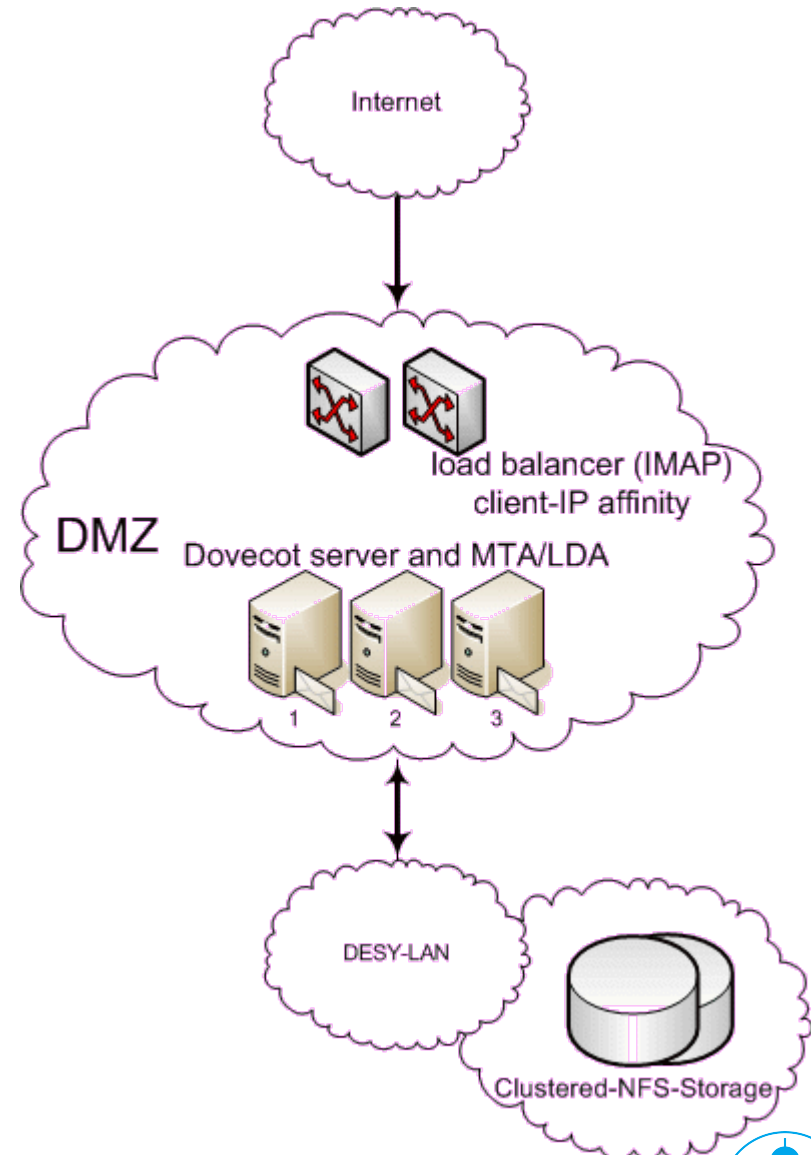
- mix several folder formats on the same server
- provide shared folders (access according to ACLs for groups of users)
- provide public folder (unlimited access after authentication)

> other folder formats possible

- mbox (indexed, faster than plain mbox files)
- compressed readonly mbox folders
- mbox with multiple mails/file (indexed, similar to UW mix, in dovecot 2.0)



- Server: 3x Sun T5120
 - UltraSPARC T2 processor with 64 simultaneous threads
 - 32GB RAM
- Storage: NFS via NetApp
 - dynamic volume sizes
 - automatic i-node allocation
 - hourly snapshots during office hours
- IMAP server: Dovecot 1.2
- Filter: Sieve-PlugIn 0.1.5
- MTA: Postfix 2.56
- MDA: Dovecot's delivery agent
- Network-Loadbalancing



- > Server: 1x Dell Poweredge 2950
 - 2x Intel® Xeon® CPU 5430 @ 2.66GHz (8cores)
 - 16GB RAM
- > Storage: locally attached
 - 8x 146GB, RAID-10 with 2 hot spares
 - 430GB netto
 - daily backups
- > IMAP server: Dovecot 1.2
- > Filter: Sieve-PlugIn 0.1.4
- > MTA: Postfix 2.33
- > MDA: Dovecot's delivery agent



> Dovecot 1.1.4 as well as 1.2rc3

- Kerberos 5 Integration via GSSAPI
 - Sun GSSAPI Stack did not work out of the box, reasons unknown
 - “While processing incoming data: Bad encryption type” and other oddities
 - MIT-Kerberos 5 (v1.6.3) did not work out of the box, either...
 - configure/autoconf has a hard coded “-lgss” for Solaris as OS
 - MIT provides “-lgssapi_krb5” instead
 - you have to get rid of all “-lgss” occurrences in the configure script yourself

> NetAPP & NFS

- creation of 100'000 files: 5m 14s
- deletion of 100'000 file: 1m 51s
- delivery of 30'000 mails: 3h, including indexing
- PINE open: 25s
 - select all for deletion: 90s
 - expunging: 30s, including indexing
- Thunderbird open: 4½m
 - move to Trash: 12m



Server side filtering

> Dovecot comes with server side filtering using Sieve language

- RFC 5228
- two implementations: CMUSieve (Cyrus) and a sieve plug-in rewritten for Dovecot
→proposed/registered **port 2000 already used** by Cisco (IP telephony)

> fairly simple syntax

- e.g.

```
require "fileinto"
if header :contains "X-Spam-Level" "*****" {
    fileinto "junkmail";
}
```

> filter rules need to be transferred to IMAP server

- homegrown techniques would mostly need root access
- use managesieve protocol (<http://tools.ietf.org/id/draft-martin-managesieve-12.txt>)
- available as a Dovecot patch/plug-in



Server side filtering from a user's perspective

- > users need easy access to server side filtering
- > in general several solutions exist
 - mixed solutions allow editing Sieve scripts and offer an interface to check, transfer, activate or deactivate the filters at the IMAP server
 - GUIs offer users support developing filters w/o having them to know the Sieve syntax
 - Sieve support integrated into mail clients offer most comfort
- > known implementations
 - Thunderbird plug-in to edit and transfer Sieve scripts (<http://sieve.mozdev.org>)
 - Mulberry has got an integrated GUI for Sieve filters
 - Squirrelmail (web based) comes with avelsieve GUI
 - websieve is a stand-alone web interface for Sieve scripts (developed for Cyrus IMAP)



Modified websieve implementation at DESY

- Web-GUI based on websieve; adapted to dovecot, Kerberos authentication added (development on-going)

[DESY homepage](#) [\[LOGOUT\]](#) [\[HELP\]](#)

[\[View/Edit Rules\]](#) [\[Add Rule\]](#) [\[Folders and ACL's\]](#) [\[Advanced Options\]](#)

Server: imap.ifh.de	Userid: friebel@IFH.DE	Used Quota: [128314 kB / 512000 kB (25.06% usage)]	Script: [default]
-------------------------------	----------------------------------	---	-----------------------------

[Save Changes](#) [Refresh](#) [Reset Values](#)

New Rule Entry for user: friebel@IFH.DE

Rule#: [1] **Priority:** 2 **Status:** ENABLED

Continue checking other rules after applying this rule
 Keep a copy in your Inbox
 Use regular expressions

IF all of **field(s):** 'from' contains
'to' contains
'subject' contains
Msg size less than (K)ilobytes
Field name contains

THEN

File Into (Mail Folder)
 Forward To (Email Address)
 Reject (Text Message)
 Discard

Hint: Use * or ? for wildcards
To invert a rule use ! as the first character of your search string

[Save Rule](#) [Clear](#)



Proof of concept: squirrelmail with avelsieve plugin

- Fairly simple PHP based webmail client
- Included in Scientific Linux5
- Built avelsieve rpm as addon
- Thought to be used as yet another sieve script generator
- Not compatible with websieve
- Available currently only in intranet (how to interpret frequent security fixes for squirrelmail?)

Server-Side Mail Filtering - Editing Mail Filtering Rule #1

Condition

The condition for the following rules is:

- Header contains
- Header contains
- Header contains

Action

Choose what to do when this rule triggers, from one of the following:

- Keep Message**
- Move to Folder** -
- Or specify a new folder: to be created under
- Redirect**
- Reject**
- Discard**
- Vacation**

Additional Actions

- STOP**
- Notify**
- Also keep copy in INBOX, marked as deleted.**
- Disable**



Migrating users from UW-IMAP to Dovecot

- > Dovecot offers some automatic conversion (mbox to Maildir)
 - advantage: all dovecot formats are supported
 - disadvantage: does not preserve UID's in mail headers
- > a migration script exists (mbox to Maildir)
 - mb2md from <http://batleth.sapientisat.org/projects/mb2md/>
 - at DESY written some glue around script for transfer of mail folders with optional conversion of only INBOX or additional local mail folders on AFS as well
 - much faster than transfer using mail client in combination with IMAP protocol
- > Other solutions do certainly exist
 - e.g. formail converts into Maildir (but not Maildir++), but does not preserve UID's



Status of Dovecot at DESY

- > Hamburg: test installation used for
 - Performance measurements
 - Configuration and failover tests
- > Zeuthen: going from pilot to production soon
 - started in Dec 2008 after very positive talk by Hildebrandt, Berlin (15000 users)
 - tests since January starting with one account
 - pilot since end of march with growing number of users (currently about 40)
 - Mix of mail clients (alpine, Thunderbird, Mac Mail, Outlook) on different platforms (Linux, Windows, MacOSX)
- > Currently no show stoppers
 - Very positive feedback from users (MUCH faster access, no other differences)
 - Minor problems were quickly solved after reporting it to the mailing list
 - No crashes, data corruption, mail loss etc. (several dovecot upgrades 1.0.6 ... 1.2.0)



Open questions and Outlook ;-)

> Need to work on websieve GUI

- Authentication and session management still buggy
- reduce complexity for casual users, e.g. offering different “levels” of knowledge (basic, advanced, expert)

> Impact of file system choice on performance

- Recent discussions suggest that XFS is superior to ext3 especially for huge folders

> impact of Maildir format when heavily used concerning (incremental) backups

- lots of checks / stats during backup necessary
- growing backup server databases
- Will be solved when mbox with multiple mails/file is in production status (end 2009?)

