

Mail server scalability

What problems do we come across?

Linear password files

On some systems, every mail delivery and pop3 connection requires a scan through the whole /etc/passwd file.

- This is a problem with many Linux distributions
- FreeBSD uses searchable databases /etc/pwd.db, /etc/spwd.db
- Don't give mail users a Unix account; have a separate user database. (Better for security, too)

Linear mbox files

If a user keeps their mail on the server, every POP3 connection requires the POP3 daemon to read the entire mail file

- Deliver each message into a separate file (Maildir), but beware you don't run out of inodes.

Too many files in one directory

- Use a hashed directory structure, e.g. /home/12/34/user

CPU limits

- Put in a faster CPU/multiple CPUs (SMP), and more RAM
- Ensure kernel parameters correctly tweaked (sockets, filehandles)
- Distribute the load across multiple boxes - clustering
- Use an efficient MTA

Disk performance

- Use softupdates (FreeBSD) or a high-performance filesystem (Linux)
- Use multiple disks, spread your mail directories across them
- Distribute the load across multiple boxes - clustering
- Enforce quotas to limit disk space used by each customer

Keep your SMTP (smarthost) and POP3 services separate

Keeping SMTP and POP3 on separate machines makes it much easier to scale your mail service.

```
pop3.example.com  -- does not relay, accepts incoming SMTP for
                  delivery to local mailboxes only
smtp.example.com  -- relays, has no local mailboxes
```

There is an additional advantage: mail routing works correctly even if one of your customers leaves (moves their domain's MX records to point somewhere else) without telling you.