



University of Cambridge Computing Service

# Exim and PCRE

## How free software hijacked my life

Philip Hazel <ph10@cus.cam.ac.uk>

- What are Exim and PCRE?
- Why did I write them?
- What were the original aims for Exim?
- How did Exim come to be generally released?
- What about PCRE?
- What have been the main problems?
- What should I have done differently?

# What are Exim and PCRE?

They are both free, open source software

Exim is a Mail Transfer Agent for Unix

- Drop-in replacement for Sendmail
- Runs on 20 operating systems

PCRE is a C/C++ regular expression library

- Supports regexs compatible with Perl 5.005
- Runs on Unix, Mac, Windows, Acorn
- Used by Exim, tin, Python, Apache, PHP, ...

# Why did I write Exim and PCRE?

**Exim:** in 1995 we were running Smail 3

I liked it, but

- 7-year old code, not standard C, much hacked
- Unwanted legacy support for UUCP addressing
- Changing requirements
  - Virtual domains
  - More checking
- Difficulty of extending the code

**PCRE:** until 1997 Exim used imported code

- Original Henry Spencer library (1986)
- Simple regular expressions
- Perl regexs are probably the most powerful
- Can't disentangle Perl's code

## What were the original aims for Exim?

- To keep Smail 3's general approach, but without the historical baggage
- To extend the capabilities
  - More controls (sender/receiver checking)
  - More reporting
  - More generality (virtual domains)
  - More flexibility (retries)
- Online systems with good connections
- 'Modern' Unixes, Standard C

# How did Exim get released?

- Exim never had a formal 'first release'
- Computer Science Dept begged an early copy
- They started running it, and told others
- We started running it on central services
- I put the source on an ftp site
- Bug reports arrived
- Requests for enhancements soon appeared
- Version went from 0.57 to 1.58 at one point
- Requests for a mailing list
- Planet Online put up a web site
- Mirrors volunteered themselves
- I started receiving T shirts...

## **What about PCRE?**

- Written summer 1997, placed on ftp site
- People found it, and started a mailing list
- There has been a trickle of enhancements

## What have been the main problems?

- A one-man band has plusses and minusses
- Documentation takes a lot of time, but is worth it
- Less knowledgeable folk are now running MTAs
- FAQ should have been done earlier
- A simple introduction is still needed
- Variations in operating systems and shells
  - Signals – default restarting or not
  - Load average finding
  - OS and architecture names
  - IP options
  - Argument types for library functions
  - Compiler pickyness
  - Things the C standard doesn't tell you
- Some OS keep changing rapidly
- Testing an MTA is difficult

## **What should I have done differently?**

- There have already been several rethinks:
  - Store handling
  - Lookup handling
  - Negative matching options
- Use less C I/O
- More radical changes to the Smail 3 conception
- Think bigger and more generally

## **Would I do it again?**

# Exim is available from

**ftp://ftp.cus.cam.ac.uk  
/pub/software/programs/exim/exim-3.xx.tar.gz**

- GNU General Public Licence
- ASCII documentation included
- PostScript (280 pages) also available
- Other formats (Texinfo, HTML, PDF) as well
- See also **<http://www.exim.org>**