

Preface

In 1995, the central servers at the University of Cambridge were running a variety of mail transfer agents, including Sendmail, Smail 3, and PP. Some years before, I had converted the systems whose mail I managed from Sendmail to Smail, to make it easier to handle the special requirements of the early 1990s in UK academic networking during the transition from a private X.25-based network to the Internet. By 1995, the transition was complete, and it was time to move on.

Up to that time, the Internet had been a pretty friendly place, and there was little need to take many precautions against hostile acts. Most sites ran open mail relays, for example. It was clear, however, that this situation was changing and that new requirements were arising. I had done some modifications to the code of Smail, but by then it was eight-year-old code, written in prestandard C, and originally designed for use in a very different environment that involved a lot of support for UUCP. I therefore decided to see if I could build a new MTA from scratch, taking the basic philosophy of Smail and extending it, but leaving out the UUCP support, which was not needed in our environment. Because I wasn't exactly sure what the outcome would be, I called it *EXperimental Internet Mailer* (Exim).

One of my colleagues in Computer Science got wind of what I was doing, begged for an evaluation copy, and promptly put it into service, even before I was running it on my hosts. He started telling others about it, so I began putting releases on an FTP site and answering email about it. The early releases were never "announced"; they just spread by word of mouth. After some time, a UK ISP volunteered to run a web site and mailing list, and it has continued to grow from there. There has been a continuous stream of comments and suggestions, and there are far more facilities in current releases than I ever planned at the start.

Although I make a point of maintaining a comprehensive reference manual, one thing that was lacking for some time was introductory and tutorial material. I kept hoping that somebody else would write something, but in the end I was asked to write a book myself. This first book described Exim 3, and was published by O'Reilly.

With the advent of Exim 4, extensive revision was required because Exim 4 differs from Exim 3 in a number of fundamental ways, and the result was a new book. Exim 4 has not stood still – many new features have appeared since this book was first published. This second edition brings it up to date for release 4.66.

Who should read this book

In these days of frequent network abuse and high volumes of junk mail, anybody who is administering a mail server needs to have a basic understanding of how

Internet mail works, and how their server processes it. This applies as much to small personal hosts as it does to large server farms.

If you are running, or thinking about running, an Exim server of any size, this book will give you a grounding in the way Exim works. In addition, it contains many examples of particular ways in which Exim can be configured to support specific mail processing requirements.

I hope this book will make life easier for those who find the reference manual difficult to work with on its own.

Changes from the previous edition of this book

One of the main reasons for revising this book was the inclusion of content-scanning facilities in Exim from release 4.50 onwards. Chapters 14 and 15 (*Message reception and policy controls* and *Rewriting addresses*) in the previous edition have been replaced by five chapters covering message reception, access control lists, address verification, scanning incoming messages, and address and header processing.

The remainder of the book has been revised to correct errors and update it for other additions to Exim, and to revise some of the general comments in the light of changing circumstances and current practice.

The previous edition contained a number of summary lists of options for routers and transports. As the book does not cover the more rarely used options, these were never complete, and became even more out-of-date as new options were added, so they are no longer present in this edition. The reference manual is the place to look for complete lists of options. Despite the removal of these lists, the body of the book is 6.5% larger than it was before.

The second appendix, which covered the regular expressions that are supported by the PCRE library, has been removed from this edition, both in the interest of reducing the book's length, and also because PCRE is now standardly installed on a number of systems and its current documentation is therefore widely available.

Organization of the book

After a short overview chapter, this book continues with a general introduction to Internet email, because this is a subject that does not seem to be well covered elsewhere. The rest of the book is devoted to explaining how Exim works, and how you can use its configuration to control what it does. Here is a detailed breakdown of the chapters:

Chapter 1 *Introduction*

This chapter is a short “executive” summary.

Chapter 2 *How Internet mail works*

This chapter is a general introduction to the way email is handled on Internet systems.

Chapter 3 *Exim overview*

This chapter contains a general overview of the way Exim works, and introduces you to the way it is configured, in particular in regard to the way messages are delivered.

Chapter 4 *Exim operations overview*

This chapter continues with more overview material, mostly about topics other than the delivery of messages.

Chapter 5 *Extending the delivery configuration*

In this chapter, we return to the subject of message delivery, and show how the configuration can be extended to support additional functionality.

Chapter 6 *Generic options that apply to all routers*

This chapter discusses the generic options that are common to all routers, which are the components of Exim that determine how a message is to be delivered.

Chapter 7 *The routers*

This chapter describes each of the routers in detail.

Chapter 8 *Generic options that apply to all transports*

This chapter discusses the generic options that are common to all transports, which are the components of Exim that actually transport messages.

Chapter 9 *The transports*

This chapter discusses each of the transports in detail.

Chapter 10 *Message filtering*

This chapter describes Exim’s filtering language that can be used both by users’ filter files and the system filter.

Chapter 11 *Exim processes*

This chapter describes the various different kinds of Exim process, and the data that they share.

Chapter 12 *Delivery errors and retrying*

This chapter is concerned with temporary delivery errors, and how Exim handles them.

Chapter 13 *Encryption, authentication, and other SMTP processing*

This chapter covers a number of topics that are concerned with the transmission and reception of messages using SMTP, including authentication and encryption.

Chapter 14 *Message reception*

This chapter contains some basic material about the way Exim receives messages.

Chapter 15 *Access control lists*

This chapter covers Access Control Lists in detail.

Chapter 16 *Address verification*

This chapter covers a number of facilities for verifying sender and recipient addresses.

Chapter 17 *Scanning incoming messages*

This chapter covers the content scanning extension to Exim, which allows the contents of messages to be scanned by external programs at SMTP time.

Chapter 18 *Address and header processing*

This chapter discusses the processing that is applied to addresses and header lines when messages are received, including address rewriting.

Chapter 19 *File and database lookups*

This is the first of three chapters that go into detail about the three main facilities that provide flexibility in Exim's configuration. They are all introduced in earlier chapters, but full details begin here.

Chapter 20 *String expansion*

This chapter covers details about Exim's string expansion mechanism.

Chapter 21 *Domain, host, and address lists*

This chapter gives more details about the several kinds of list that can appear in Exim configurations.

Chapter 22 *Miscellany*

This chapter collects together a number of items that do not fit naturally into the other chapters, but which are too small to warrant individual chapters of their own.

Chapter 23 *Command-line interface to Exim*

This chapter describes the options and arguments that are used to control what a call to Exim actually does.

Chapter 24 *Administering Exim*

This chapter discusses a number of topics concerned with administration, and describes the utility programs that are available to help with this.

Chapter 25 *Building and installing Exim*

This chapter describes how to build and install Exim from the source distribution.

Appendix *Summary of string expansion*

The appendix contains a summary of string expansion items.

Conventions used in this book

The following is a list of the typographical conventions used in this book:

Italic

Used for file and directory names, program and command names, host and domain names, email addresses, mail headers, URLs, and new terms.

****Bold****

Used for the names of Exim routers, transports, and authenticators.

Slanted

Used for the names of Exim variables.

Constant Width

Used in examples to show the contents of files or the output from commands, and in the text to mark Exim options or other strings that appear literally in configuration files.

<Constant Italic>

Used to indicate variable options, keywords, or text that the user is to replace with an actual value.

****Constant Bold****

Used in examples to show commands or other text that should be typed literally by the user.

Cross-references within the book are by section number, and they are often shown in parentheses like this: (☞ 7.2).

Suggestions and comments on this book

The author and publisher welcome feedback from all readers. If you have any comments on this book, would like to make a suggestion, or have noticed an error, please email us on the appropriate address:

comments@uit.co.uk
suggestions@uit.co.uk
errors@uit.co.uk

Acknowledgments

I could not have produced Exim without the support and assistance of many people and organizations, and in particular, my colleagues in the Computing Service at the University of Cambridge. The management allowed me to write Exim, and once it appeared, the Computing Service has supported its use around the University and elsewhere.

Since the first release, many people have sent suggestions and code for improvements or new features, and fixes for problems, and this continues unabated, so much so that there are now too many contributors to list here individually. Their names are, however, included in an *ACKNOWLEDGEMENTS* file in Exim distributions.

While writing the different versions of this book, I have continued to enjoy the support of my colleagues and the Exim community. My wife Judith was not only generally supportive, but also read both an early draft and a later version as a professional copy editor, and found many places where I was unclear or inconsistent. Ken Bailey made some useful comments about some of the early chapters. John Horne read an early draft of the Exim 3 book and made suggestions that helped me to put the material into a more accessible order, and then read it again in a late draft, providing further useful feedback. John also read and commented on the entirely new chapter on access control for the first edition of the Exim 4 book. Michael Shappe read the early chapters and provided me with yet more useful technical and editorial feedback. David McLaughlin read this new edition and sent helpful comments.

My editor at O'Reilly for the Exim 3 book was Andy Oram. His comments and guidance had a great effect on the form and shape of the finished book, and much of what he did has carried over into the Exim 4 book.

Last but by no means least, I must thank Niall Mansfield, of UIT Cambridge, who came to my rescue when I needed a new publisher.