

# **Specification of the Exim Mail Transfer Agent**

**Exim Maintainers**

# **Specification of the Exim Mail Transfer Agent**

Author: Exim Maintainers

Copyright © 2017 University of Cambridge

Revision 4.90 13 Dec 2017

# Contents

<b>1. Introduction .....</b>	<b>1</b>
1.1 Exim documentation .....	1
1.2 FTP and web sites .....	2
1.3 Mailing lists .....	2
1.4 Bug reports .....	3
1.5 Where to find the Exim distribution .....	3
1.6 Limitations .....	3
1.7 Run time configuration .....	4
1.8 Calling interface .....	4
1.9 Terminology .....	4
<b>2. Incorporated code .....</b>	<b>6</b>
<b>3. How Exim receives and delivers mail .....</b>	<b>8</b>
3.1 Overall philosophy .....	8
3.2 Policy control .....	8
3.3 User filters .....	8
3.4 Message identification .....	9
3.5 Receiving mail .....	9
3.6 Handling an incoming message .....	10
3.7 Life of a message .....	10
3.8 Processing an address for delivery .....	11
3.9 Processing an address for verification .....	12
3.10 Running an individual router .....	12
3.11 Duplicate addresses .....	13
3.12 Router preconditions .....	13
3.13 Delivery in detail .....	14
3.14 Retry mechanism .....	15
3.15 Temporary delivery failure .....	15
3.16 Permanent delivery failure .....	15
3.17 Failures to deliver bounce messages .....	16
<b>4. Building and installing Exim .....</b>	<b>17</b>
4.1 Unpacking .....	17
4.2 Multiple machine architectures and operating systems .....	17
4.3 PCRE library .....	17
4.4 DBM libraries .....	17
4.5 Pre-building configuration .....	18
4.6 Support for iconv() .....	19
4.7 Including TLS/SSL encryption support .....	19
4.8 Use of tcpwrappers .....	20
4.9 Including support for IPv6 .....	20
4.10 Dynamically loaded lookup module support .....	21
4.11 The building process .....	21
4.12 Output from “make” .....	21
4.13 Overriding build-time options for Exim .....	21
4.14 OS-specific header files .....	23
4.15 Overriding build-time options for the monitor .....	23
4.16 Installing Exim binaries and scripts .....	24
4.17 Installing info documentation .....	25
4.18 Setting up the spool directory .....	25
4.19 Testing .....	25

4.20	Replacing another MTA with Exim .....	26
4.21	Upgrading Exim .....	27
4.22	Stopping the Exim daemon on Solaris .....	27
<b>5.</b>	<b>The Exim command line .....</b>	<b>28</b>
5.1	Setting options by program name .....	28
5.2	Trusted and admin users .....	28
5.3	Command line options .....	29
<b>6.</b>	<b>The Exim run time configuration file .....</b>	<b>51</b>
6.1	Using a different configuration file .....	51
6.2	Configuration file format .....	52
6.3	File inclusions in the configuration file .....	53
6.4	Macros in the configuration file .....	53
6.5	Macro substitution .....	53
6.6	Redefining macros .....	54
6.7	Overriding macro values .....	54
6.8	Example of macro usage .....	54
6.9	Builtin macros .....	54
6.10	Conditional skips in the configuration file .....	55
6.11	Common option syntax .....	55
6.12	Boolean options .....	55
6.13	Integer values .....	56
6.14	Octal integer values .....	56
6.15	Fixed point numbers .....	56
6.16	Time intervals .....	56
6.17	String values .....	56
6.18	Expanded strings .....	57
6.19	User and group names .....	57
6.20	List construction .....	57
6.21	Changing list separators .....	57
6.22	Empty items in lists .....	58
6.23	Format of driver configurations .....	58
<b>7.</b>	<b>The default configuration file .....</b>	<b>60</b>
7.1	Main configuration settings .....	60
7.2	ACL configuration .....	63
7.3	Router configuration .....	66
7.4	Transport configuration .....	68
7.5	Default retry rule .....	69
7.6	Rewriting configuration .....	69
7.7	Authenticators configuration .....	69
<b>8.</b>	<b>Regular expressions .....</b>	<b>71</b>
<b>9.</b>	<b>File and database lookups .....</b>	<b>72</b>
9.1	Examples of different lookup syntax .....	72
9.2	Lookup types .....	73
9.3	Single-key lookup types .....	73
9.4	Query-style lookup types .....	75
9.5	Temporary errors in lookups .....	76
9.6	Default values in single-key lookups .....	76
9.7	Partial matching in single-key lookups .....	77
9.8	Lookup caching .....	78

9.9	Quoting lookup data .....	78
9.10	More about dnsdb .....	79
9.11	Dnsdb lookup modifiers .....	79
9.12	Pseudo dnsdb record types .....	80
9.13	Multiple dnsdb lookups .....	81
9.14	More about LDAP .....	81
9.15	Format of LDAP queries .....	81
9.16	LDAP quoting .....	82
9.17	LDAP connections .....	82
9.18	LDAP authentication and control information .....	83
9.19	Format of data returned by LDAP .....	85
9.20	More about NIS+ .....	85
9.21	SQL lookups .....	86
9.22	More about MySQL, PostgreSQL, Oracle, InterBase, and Redis .....	86
9.23	Specifying the server in the query .....	87
9.24	Special MySQL features .....	87
9.25	Special PostgreSQL features .....	88
9.26	More about SQLite .....	88
9.27	More about Redis .....	88
<b>10.</b>	<b>Domain, host, address, and local part lists .....</b>	<b>89</b>
10.1	Expansion of lists .....	89
10.2	Negated items in lists .....	89
10.3	File names in lists .....	90
10.4	An lsearch file is not an out-of-line list .....	90
10.5	Named lists .....	90
10.6	Named lists compared with macros .....	91
10.7	Named list caching .....	91
10.8	Domain lists .....	92
10.9	Host lists .....	94
10.10	Special host list patterns .....	94
10.11	Host list patterns that match by IP address .....	94
10.12	Host list patterns for single-key lookups by host address .....	95
10.13	Host list patterns that match by host name .....	96
10.14	Behaviour when an IP address or name cannot be found .....	97
10.15	Mixing wildcarded host names and addresses in host lists .....	97
10.16	Temporary DNS errors when looking up host information .....	98
10.17	Host list patterns for single-key lookups by host name .....	98
10.18	Host list patterns for query-style lookups .....	98
10.19	Address lists .....	99
10.20	Case of letters in address lists .....	101
10.21	Local part lists .....	101
<b>11.</b>	<b>String expansions .....</b>	<b>102</b>
11.1	Literal text in expanded strings .....	102
11.2	Character escape sequences in expanded strings .....	102
11.3	Testing string expansions .....	102
11.4	Forced expansion failure .....	103
11.5	Expansion items .....	103
11.6	Expansion operators .....	114
11.7	Expansion conditions .....	120
11.8	Combining expansion conditions .....	127
11.9	Expansion variables .....	128
<b>12.</b>	<b>Embedded Perl .....</b>	<b>147</b>

12.1	Setting up so Perl can be used .....	147
12.2	Calling Perl subroutines .....	147
12.3	Calling Exim functions from Perl .....	148
12.4	Use of standard output and error by Perl .....	148
<b>13.</b>	<b>Starting the daemon and the use of network interfaces .....</b>	<b>149</b>
13.1	Starting a listening daemon .....	149
13.2	Special IP listening addresses .....	150
13.3	Overriding local_interfaces and daemon_smtp_ports .....	150
13.4	Support for the obsolete SSMTP (or SMTPS) protocol .....	150
13.5	IPv6 address scopes .....	151
13.6	Disabling IPv6 .....	151
13.7	Examples of starting a listening daemon .....	151
13.8	Recognizing the local host .....	152
13.9	Delivering to a remote host .....	152
<b>14.</b>	<b>Main configuration .....</b>	<b>153</b>
14.1	Miscellaneous .....	153
14.2	Exim parameters .....	153
14.3	Privilege controls .....	153
14.4	Logging .....	154
14.5	Frozen messages .....	154
14.6	Data lookups .....	154
14.7	Message ids .....	154
14.8	Embedded Perl Startup .....	154
14.9	Daemon .....	154
14.10	Resource control .....	155
14.11	Policy controls .....	155
14.12	Callout cache .....	156
14.13	TLS .....	156
14.14	Local user handling .....	156
14.15	All incoming messages (SMTP and non-SMTP) .....	157
14.16	Non-SMTP incoming messages .....	157
14.17	Incoming SMTP messages .....	157
14.18	SMTP extensions .....	157
14.19	Processing messages .....	158
14.20	System filter .....	158
14.21	Routing and delivery .....	158
14.22	Bounce and warning messages .....	159
14.23	Alphabetical list of main options .....	159
<b>15.</b>	<b>Generic options for routers .....</b>	<b>207</b>
<b>16.</b>	<b>The accept router .....</b>	<b>221</b>
<b>17.</b>	<b>The dnslookup router .....</b>	<b>222</b>
17.1	Problems with DNS lookups .....	222
17.2	Declining addresses by dnslookup .....	222
17.3	Private options for dnslookup .....	223
17.4	Effect of qualify_single and search_parents .....	225
<b>18.</b>	<b>The ipliteral router .....</b>	<b>226</b>
<b>19.</b>	<b>The iplookup router .....</b>	<b>227</b>

<b>20. The manualroute router .....</b>	<b>229</b>
20.1 Private options for manualroute .....	229
20.2 Routing rules in route_list .....	230
20.3 Routing rules in route_data .....	231
20.4 Format of the list of hosts .....	231
20.5 Format of one host item .....	232
20.6 How the list of hosts is used .....	232
20.7 How the options are used .....	233
20.8 Manualroute examples .....	233
<b>21. The queryprogram router .....</b>	<b>236</b>
<b>22. The redirect router .....</b>	<b>238</b>
22.1 Redirection data .....	238
22.2 Forward files and address verification .....	238
22.3 Interpreting redirection data .....	239
22.4 Items in a non-filter redirection list .....	239
22.5 Redirecting to a local mailbox .....	239
22.6 Special items in redirection lists .....	240
22.7 Duplicate addresses .....	242
22.8 Repeated redirection expansion .....	242
22.9 Errors in redirection lists .....	242
22.10 Private options for the redirect router .....	242
<b>23. Environment for running local transports .....</b>	<b>251</b>
23.1 Concurrent deliveries .....	251
23.2 Uids and gids .....	251
23.3 Current and home directories .....	252
23.4 Expansion variables derived from the address .....	252
<b>24. Generic options for transports .....</b>	<b>253</b>
<b>25. Address batching in local transports .....</b>	<b>260</b>
<b>26. The appendfile transport .....</b>	<b>262</b>
26.1 The file and directory options .....	262
26.2 Private options for appendfile .....	263
26.3 Operational details for appending .....	272
26.4 Operational details for delivery to a new file .....	274
26.5 Maildir delivery .....	274
26.6 Using tags to record message sizes .....	275
26.7 Using a maildirsized file .....	276
26.8 Mailstore delivery .....	276
26.9 Non-special new file delivery .....	276
<b>27. The autoreply transport .....</b>	<b>277</b>
27.1 Private options for autoreply .....	277
<b>28. The lmtp transport .....</b>	<b>280</b>
<b>29. The pipe transport .....</b>	<b>282</b>
29.1 Concurrent delivery .....	282

29.2	Returned status and data .....	282
29.3	How the command is run .....	283
29.4	Environment variables .....	284
29.5	Private options for pipe .....	284
29.6	Using an external local delivery agent .....	289
<b>30.</b>	<b>The smtp transport .....</b>	<b>291</b>
30.1	Multiple messages on a single connection .....	291
30.2	Use of the \$host and \$host_address variables .....	291
30.3	Use of \$tls_cipher and \$tls_peerdn .....	291
30.4	Private options for smtp .....	291
30.5	How the limits for the number of hosts to try are used .....	302
<b>31.</b>	<b>Address rewriting .....</b>	<b>304</b>
31.1	Explicitly configured address rewriting .....	304
31.2	When does rewriting happen? .....	304
31.3	Testing the rewriting rules that apply on input .....	305
31.4	Rewriting rules .....	305
31.5	Rewriting patterns .....	306
31.6	Rewriting replacements .....	307
31.7	Rewriting flags .....	307
31.8	Flags specifying which headers and envelope addresses to rewrite .....	307
31.9	The SMTP-time rewriting flag .....	307
31.10	Flags controlling the rewriting process .....	308
31.11	Rewriting examples .....	308
<b>32.</b>	<b>Retry configuration .....</b>	<b>310</b>
32.1	Changing retry rules .....	310
32.2	Format of retry rules .....	310
32.3	Choosing which retry rule to use for address errors .....	311
32.4	Choosing which retry rule to use for host and message errors .....	311
32.5	Retry rules for specific errors .....	312
32.6	Retry rules for specified senders .....	313
32.7	Retry parameters .....	314
32.8	Retry rule examples .....	314
32.9	Timeout of retry data .....	315
32.10	Long-term failures .....	315
32.11	Deliveries that work intermittently .....	316
<b>33.</b>	<b>SMTP authentication .....</b>	<b>317</b>
33.1	Generic options for authenticators .....	318
33.2	The AUTH parameter on MAIL commands .....	320
33.3	Authentication on an Exim server .....	320
33.4	Testing server authentication .....	321
33.5	Authentication by an Exim client .....	322
<b>34.</b>	<b>The plaintext authenticator .....</b>	<b>323</b>
34.1	Plaintext options .....	323
34.2	Using plaintext in a server .....	323
34.3	The PLAIN authentication mechanism .....	323
34.4	The LOGIN authentication mechanism .....	324
34.5	Support for different kinds of authentication .....	325
34.6	Using plaintext in a client .....	325



<b>35. The cram_md5 authenticator .....</b>	<b>327</b>
35.1 Using cram_md5 as a server .....	327
35.2 Using cram_md5 as a client .....	327
<b>36. The cyrus_sasl authenticator .....</b>	<b>329</b>
36.1 Using cyrus_sasl as a server .....	329
<b>37. The dovecot authenticator .....</b>	<b>331</b>
<b>38. The gssapi authenticator .....</b>	<b>332</b>
38.1 gssapi auth variables .....	333
<b>39. The heimdal_gssapi authenticator .....</b>	<b>334</b>
39.1 heimdal_gssapi auth variables .....	334
<b>40. The spa authenticator .....</b>	<b>335</b>
40.1 Using spa as a server .....	335
40.2 Using spa as a client .....	335
<b>41. The tls authenticator .....</b>	<b>337</b>
<b>42. Encrypted SMTP connections using TLS/SSL .....</b>	<b>338</b>
42.1 Support for the legacy “ssmtp” (aka “smtps”) protocol .....	338
42.2 OpenSSL vs GnuTLS .....	338
42.3 GnuTLS parameter computation .....	339
42.4 Requiring specific ciphers in OpenSSL .....	340
42.5 Requiring specific ciphers or other parameters in GnuTLS .....	341
42.6 Configuring an Exim server to use TLS .....	341
42.7 Requesting and verifying client certificates .....	343
42.8 Revoked certificates .....	343
42.9 Configuring an Exim client to use TLS .....	344
42.10 Use of TLS Server Name Indication .....	345
42.11 Multiple messages on the same encrypted TCP/IP connection .....	346
42.12 Certificates and all that .....	347
42.13 Certificate chains .....	347
42.14 Self-signed certificates .....	347
<b>43. Access control lists .....</b>	<b>349</b>
43.1 Testing ACLs .....	349
43.2 Specifying when ACLs are used .....	349
43.3 The non-SMTP ACLs .....	350
43.4 The SMTP connect ACL .....	350
43.5 The EHLO/HELO ACL .....	350
43.6 The DATA ACLs .....	350
43.7 The SMTP DKIM ACL .....	351
43.8 The SMTP MIME ACL .....	351
43.9 The SMTP PRDR ACL .....	351
43.10 The QUIT ACL .....	351
43.11 The not-QUIT ACL .....	352
43.12 Finding an ACL to use .....	352

43.13	ACL return codes .....	353
43.14	Unset ACL options .....	353
43.15	Data for message ACLs .....	354
43.16	Data for non-message ACLs .....	354
43.17	Format of an ACL .....	354
43.18	ACL verbs .....	355
43.19	ACL variables .....	356
43.20	Condition and modifier processing .....	357
43.21	ACL modifiers .....	358
43.22	Use of the control modifier .....	362
43.23	Summary of message fixup control .....	366
43.24	Adding header lines in ACLs .....	366
43.25	Removing header lines in ACLs .....	367
43.26	ACL conditions .....	368
43.27	Using DNS lists .....	372
43.28	Specifying the IP address for a DNS list lookup .....	373
43.29	DNS lists keyed on domain names .....	373
43.30	Multiple explicit keys for a DNS list .....	374
43.31	Data returned by DNS lists .....	374
43.32	Variables set from DNS lists .....	375
43.33	Additional matching conditions for DNS lists .....	375
43.34	Negated DNS matching conditions .....	376
43.35	Handling multiple DNS records from a DNS list .....	376
43.36	Detailed information from merged DNS lists .....	377
43.37	DNS lists and IPv6 .....	378
43.38	Rate limiting incoming messages .....	378
43.39	Ratelimit options for what is being measured .....	379
43.40	Ratelimit update modes .....	380
43.41	Ratelimit options for handling fast clients .....	380
43.42	Limiting the rate of different events .....	381
43.43	Using rate limiting .....	381
43.44	Address verification .....	382
43.45	Callout verification .....	383
43.46	Additional parameters for callouts .....	384
43.47	Callout caching .....	386
43.48	Sender address verification reporting .....	386
43.49	Redirection while verifying .....	387
43.50	Client SMTP authorization (CSA) .....	387
43.51	Bounce address tag validation .....	388
43.52	Using an ACL to control relaying .....	389
43.53	Checking a relay configuration .....	390
<b>44.</b>	<b>Content scanning at ACL time .....</b>	<b>391</b>
44.1	Scanning for viruses .....	391
44.2	Scanning with SpamAssassin and Rspamd .....	395
44.3	Calling SpamAssassin from an Exim ACL .....	397
44.4	Scanning MIME parts .....	398
44.5	Scanning with regular expressions .....	401
<b>45.</b>	<b>Adding a local scan function to Exim .....</b>	<b>402</b>
45.1	Building Exim to use a local scan function .....	402
45.2	API for local_scan() .....	402
45.3	Configuration options for local_scan() .....	403
45.4	Available Exim variables .....	404
45.5	Structure of header lines .....	406
45.6	Structure of recipient items .....	406

45.7 Available Exim functions .....	407
45.8 More about Exim's memory handling .....	411
<b>46. System-wide message filtering .....</b>	<b>412</b>
46.1 Specifying a system filter .....	412
46.2 Testing a system filter .....	412
46.3 Contents of a system filter .....	412
46.4 Additional variable for system filters .....	413
46.5 Defer, freeze, and fail commands for system filters .....	413
46.6 Adding and removing headers in a system filter .....	414
46.7 Setting an errors address in a system filter .....	414
46.8 Per-address filtering .....	415
<b>47. Message processing .....</b>	<b>416</b>
47.1 Submission mode for non-local messages .....	416
47.2 Line endings .....	417
47.3 Unqualified addresses .....	417
47.4 The UUCP From line .....	418
47.5 Resent- header lines .....	418
47.6 The Auto-Submitted: header line .....	419
47.7 The Bcc: header line .....	419
47.8 The Date: header line .....	419
47.9 The Delivery-date: header line .....	419
47.10 The Envelope-to: header line .....	419
47.11 The From: header line .....	419
47.12 The Message-ID: header line .....	420
47.13 The Received: header line .....	420
47.14 The References: header line .....	420
47.15 The Return-path: header line .....	420
47.16 The Sender: header line .....	420
47.17 Adding and removing header lines in routers and transports .....	421
47.18 Constructed addresses .....	422
47.19 Case of local parts .....	423
47.20 Dots in local parts .....	423
47.21 Rewriting addresses .....	423
<b>48. SMTP processing .....</b>	<b>424</b>
48.1 Outgoing SMTP and LMTP over TCP/IP .....	424
48.2 Errors in outgoing SMTP .....	425
48.3 Incoming SMTP messages over TCP/IP .....	426
48.4 Unrecognized SMTP commands .....	428
48.5 Syntax and protocol errors in SMTP commands .....	428
48.6 Use of non-mail SMTP commands .....	428
48.7 The VRFY and EXPN commands .....	428
48.8 The ETRN command .....	428
48.9 Incoming local SMTP .....	429
48.10 Outgoing batched SMTP .....	429
48.11 Incoming batched SMTP .....	430
<b>49. Customizing bounce and warning messages .....</b>	<b>431</b>
49.1 Customizing bounce messages .....	431
49.2 Customizing warning messages .....	432
<b>50. Some common configuration settings .....</b>	<b>433</b>

50.1	Sending mail to a smart host .....	433
50.2	Using Exim to handle mailing lists .....	433
50.3	Syntax errors in mailing lists .....	433
50.4	Re-expansion of mailing lists .....	434
50.5	Closed mailing lists .....	434
50.6	Variable Envelope Return Paths (VERP) .....	435
50.7	Virtual domains .....	436
50.8	Multiple user mailboxes .....	437
50.9	Simplified vacation processing .....	438
50.10	Taking copies of mail .....	438
50.11	Intermittently connected hosts .....	438
50.12	Exim on the upstream server host .....	438
50.13	Exim on the intermittently connected client host .....	439
<b>51.</b>	<b>Using Exim as a non-queueing client .....</b>	<b>440</b>
<b>52.</b>	<b>Log files .....</b>	<b>442</b>
52.1	Where the logs are written .....	442
52.2	Logging to local files that are periodically “cycled” .....	443
52.3	Datestamped log files .....	443
52.4	Logging to syslog .....	444
52.5	Log line flags .....	445
52.6	Logging message reception .....	445
52.7	Logging deliveries .....	446
52.8	Discarded deliveries .....	447
52.9	Deferred deliveries .....	447
52.10	Delivery failures .....	447
52.11	Fake deliveries .....	448
52.12	Completion .....	448
52.13	Summary of Fields in Log Lines .....	448
52.14	Other log entries .....	449
52.15	Reducing or increasing what is logged .....	449
52.16	Message log .....	453
<b>53.</b>	<b>Exim utilities .....</b>	<b>455</b>
53.1	Finding out what Exim processes are doing (exiwhat) .....	455
53.2	Selective queue listing (exiqgrep) .....	455
53.3	Summarizing the queue (exiqsumm) .....	456
53.4	Extracting specific information from the log (exigrep) .....	457
53.5	Selecting messages by various criteria (exipick) .....	457
53.6	Cycling log files (exicyclog) .....	458
53.7	Mail statistics (eximstats) .....	458
53.8	Checking access policy (exim_checkaccess) .....	459
53.9	Making DBM files (exim_dbmbuild) .....	459
53.10	Finding individual retry times (exinext) .....	460
53.11	Hints database maintenance .....	460
53.12	exim_dumpdb .....	461
53.13	exim_tidydb .....	461
53.14	exim_fixdb .....	462
53.15	Mailbox maintenance (exim_lock) .....	462
<b>54.</b>	<b>The Exim monitor .....</b>	<b>464</b>
54.1	Running the monitor .....	464
54.2	The stripcharts .....	464

54.3	Main action buttons .....	465
54.4	The log display .....	465
54.5	The queue display .....	466
54.6	The queue menu .....	466
<b>55.</b>	<b>Security considerations .....</b>	<b>469</b>
55.1	Building a more “hardened” Exim .....	469
55.2	Root privilege .....	469
55.3	Running Exim without privilege .....	471
55.4	Delivering to local files .....	472
55.5	Running local commands .....	472
55.6	Trust in configuration data .....	472
55.7	IPv4 source routing .....	473
55.8	The VRFY, EXPN, and ETRN commands in SMTP .....	473
55.9	Privileged users .....	473
55.10	Spool files .....	473
55.11	Use of argv[0] .....	474
55.12	Use of %f formatting .....	474
55.13	Embedded Exim path .....	474
55.14	Dynamic module directory .....	474
55.15	Use of sprintf() .....	474
55.16	Use of debug_printf() and log_write() .....	474
55.17	Use of strcat() and strcpy() .....	474
<b>56.</b>	<b>Format of spool files .....</b>	<b>475</b>
56.1	Format of the -H file .....	475
56.2	Format of the -D file .....	479
<b>57.</b>	<b>Support for DKIM (DomainKeys Identified Mail) .....</b>	<b>480</b>
57.1	Signing outgoing messages .....	480
57.2	Verifying DKIM signatures in incoming mail .....	481
<b>58.</b>	<b>Proxies .....</b>	<b>485</b>
58.1	Inbound proxies .....	485
58.2	Outbound proxies .....	485
58.3	Logging .....	486
<b>59.</b>	<b>Internationalisation .....</b>	<b>487</b>
59.1	MTA operations .....	487
59.2	MDA operations .....	487
<b>60.</b>	<b>Events .....</b>	<b>489</b>
<b>61.</b>	<b>Adding new drivers or lookup types .....</b>	<b>491</b>
	<b>Options index .....</b>	<b>492</b>
	<b>Variables index .....</b>	<b>498</b>
	<b>Concept index .....</b>	<b>500</b>



# 1. Introduction

Exim is a mail transfer agent (MTA) for hosts that are running Unix or Unix-like operating systems. It was designed on the assumption that it would be run on hosts that are permanently connected to the Internet. However, it can be used on intermittently connected hosts with suitable configuration adjustments.

Configuration files currently exist for the following operating systems: AIX, BSD/OS (aka BSDI), Darwin (Mac OS X), DGUX, Dragonfly, FreeBSD, GNU/Hurd, GNU/Linux, HI-OSF (Hitachi), HI-UX, HP-UX, IRIX, MIPS RISCOS, NetBSD, OpenBSD, OpenUNIX, QNX, SCO, SCO SVR4.2 (aka UNIX-SV), Solaris (aka SunOS5), SunOS4, Tru64-Unix (formerly Digital UNIX, formerly DEC-OSF1), Ultrix, and Unixware. Some of these operating systems are no longer current and cannot easily be tested, so the configuration files may no longer work in practice.

There are also configuration files for compiling Exim in the Cygwin environment that can be installed on systems running Windows. However, this document does not contain any information about running Exim in the Cygwin environment.

The terms and conditions for the use and distribution of Exim are contained in the file *NOTICE*. Exim is distributed under the terms of the GNU General Public Licence, a copy of which may be found in the file *LICENCE*.

The use, supply or promotion of Exim for the purpose of sending bulk, unsolicited electronic mail is incompatible with the basic aims of the program, which revolve around the free provision of a service that enhances the quality of personal communications. The author of Exim regards indiscriminate mass-mailing as an antisocial, irresponsible abuse of the Internet.

Exim owes a great deal to Smail 3 and its author, Ron Karr. Without the experience of running and working on the Smail 3 code, I could never have contemplated starting to write a new MTA. Many of the ideas and user interfaces were originally taken from Smail 3, though the actual code of Exim is entirely new, and has developed far beyond the initial concept.

Many people, both in Cambridge and around the world, have contributed to the development and the testing of Exim, and to porting it to various operating systems. I am grateful to them all. The distribution now contains a file called *ACKNOWLEDGMENTS*, in which I have started recording the names of contributors.

## 1.1 Exim documentation

This edition of the Exim specification applies to version 4.90 of Exim. Substantive changes from the 4.89 edition are marked in some renditions of the document; this paragraph is so marked if the rendition is capable of showing a change indicator.

This document is very much a reference manual; it is not a tutorial. The reader is expected to have some familiarity with the SMTP mail transfer protocol and with general Unix system administration. Although there are some discussions and examples in places, the information is mostly organized in a way that makes it easy to look up, rather than in a natural order for sequential reading. Furthermore, the manual aims to cover every aspect of Exim in detail, including a number of rarely-used, special-purpose features that are unlikely to be of very wide interest.

An “easier” discussion of Exim which provides more in-depth explanatory, introductory, and tutorial material can be found in a book entitled *The Exim SMTP Mail Server* (second edition, 2007), published by UIT Cambridge (<http://www.uit.co.uk/exim-book/>).

This book also contains a chapter that gives a general introduction to SMTP and Internet mail. Inevitably, however, the book is unlikely to be fully up-to-date with the latest release of Exim. (Note that the earlier book about Exim, published by O'Reilly, covers Exim 3, and many things have changed in Exim 4.)

If you are using a Debian distribution of Exim, you will find information about Debian-specific features in the file `/usr/share/doc/exim4-base/README.Debian`. The command `man update-exim.conf` is another source of Debian-specific information.

As the program develops, there may be features in newer versions that have not yet made it into this document, which is updated only when the most significant digit of the fractional part of the version number changes. Specifications of new features that are not yet in this manual are placed in the file *doc/NewStuff* in the Exim distribution.

Some features may be classified as “experimental”. These may change incompatibly while they are developing, or even be withdrawn. For this reason, they are not documented in this manual. Information about experimental features can be found in the file *doc/experimental.txt*.

All changes to the program (whether new features, bug fixes, or other kinds of change) are noted briefly in the file called *doc/ChangeLog*.

This specification itself is available as an ASCII file in *doc/spec.txt* so that it can easily be searched with a text editor. Other files in the *doc* directory are:

<i>OptionLists.txt</i>	list of all options in alphabetical order
<i>dbm.discuss.txt</i>	discussion about DBM libraries
<i>exim.8</i>	a man page of Exim’s command line options
<i>experimental.txt</i>	documentation of experimental features
<i>filter.txt</i>	specification of the filter language
<i>Exim3.upgrade</i>	upgrade notes from release 2 to release 3
<i>Exim4.upgrade</i>	upgrade notes from release 3 to release 4
<i>openssl.txt</i>	installing a current OpenSSL release

The main specification and the specification of the filtering language are also available in other formats (HTML, PostScript, PDF, and Texinfo). Section 1.5 below tells you how to get hold of these.

## 1.2 FTP and web sites

The primary site for Exim source distributions is currently the University of Cambridge’s FTP site, whose contents are described in *Where to find the Exim distribution* below. In addition, there is a web site and an FTP site at **exim.org**. These are now also hosted at the University of Cambridge. The **exim.org** site was previously hosted for a number of years by Energis Squared, formerly Planet Online Ltd, whose support I gratefully acknowledge.

As well as Exim distribution tar files, the Exim web site contains a number of differently formatted versions of the documentation. A recent addition to the online information is the Exim wiki (<http://wiki.exim.org>), which contains what used to be a separate FAQ, as well as various other examples, tips, and know-how that have been contributed by Exim users.

An Exim Bugzilla exists at <https://bugs.exim.org>. You can use this to report bugs, and also to add items to the wish list. Please search first to check that you are not duplicating a previous entry.

## 1.3 Mailing lists

The following Exim mailing lists exist:

<i>exim-announce@exim.org</i>	Moderated, low volume announcements list
<i>exim-users@exim.org</i>	General discussion list
<i>exim-dev@exim.org</i>	Discussion of bugs, enhancements, etc.
<i>exim-cvs@exim.org</i>	Automated commit messages from the VCS

You can subscribe to these lists, change your existing subscriptions, and view or search the archives via the mailing lists link on the Exim home page. If you are using a Debian distribution of Exim, you may wish to subscribe to the Debian-specific mailing list *pkg-exim4-users@lists.alioth.debian.org* via this web page:

**<http://lists.alioth.debian.org/mailman/listinfo/pkg-exim4-users>**

Please ask Debian-specific questions on this list and not on the general Exim lists.



## 1.4 Bug reports

Reports of obvious bugs can be emailed to [bugs@exim.org](mailto:bugs@exim.org) or reported via the Bugzilla (<https://bugs.exim.org>). However, if you are unsure whether some behaviour is a bug or not, the best thing to do is to post a message to the *exim-dev* mailing list and have it discussed.

## 1.5 Where to find the Exim distribution

The master ftp site for the Exim distribution is

**<ftp://ftp.exim.org/pub/exim>**

The file references that follow are relative to the *exim* directories at these sites. There are now quite a number of independent mirror sites around the world. Those that I know about are listed in the file called *Mirrors*.

Within the *exim* directory there are subdirectories called *exim3* (for previous Exim 3 distributions), *exim4* (for the latest Exim 4 distributions), and *Testing* for testing versions. In the *exim4* subdirectory, the current release can always be found in files called

*exim-n.nn.tar.gz*  
*exim-n.nn.tar.bz2*

where *n.nn* is the highest such version number in the directory. The two files contain identical data; the only difference is the type of compression. The *.bz2* file is usually a lot smaller than the *.gz* file.

The distributions will be PGP signed by an individual key of the Release Coordinator. This key will have a uid containing an email address in the *exim.org* domain and will have signatures from other people, including other Exim maintainers. We expect that the key will be in the "strong set" of PGP keys. There should be a trust path to that key from Nigel Metheringham's PGP key, a version of which can be found in the release directory in the file *nigel-pubkey.asc*. All keys used will be available in public keyserver pools, such as *pool.sks-keyservers.net*.

At time of last update, releases were being made by Phil Pennock and signed with key *0x403043153903637F*, although that key is expected to be replaced in 2013. A trust path from Nigel's key to Phil's can be observed at <https://www.security.spodhuis.org/exim-trustpath>.

Releases have also been authorized to be performed by Todd Lyons who signs with key *0xC4F4F94804D29EBA*. A direct trust path exists between previous RE Phil Pennock and Todd Lyons through a common associate.

The signatures for the tar bundles are in:

*exim-n.nn.tar.gz.asc*  
*exim-n.nn.tar.bz2.asc*

For each released version, the log of changes is made separately available in a separate file in the directory *ChangeLogs* so that it is possible to find out what has changed without having to download the entire distribution.

The main distribution contains ASCII versions of this specification and other documentation; other formats of the documents are available in separate files inside the *exim4* directory of the FTP site:

*exim-html-n.nn.tar.gz*  
*exim-pdf-n.nn.tar.gz*  
*exim-postscript-n.nn.tar.gz*  
*exim-texinfo-n.nn.tar.gz*

These tar files contain only the *doc* directory, not the complete distribution, and are also available in *.bz2* as well as *.gz* forms.

## 1.6 Limitations

- Exim is designed for use as an Internet MTA, and therefore handles addresses in RFC 2822 domain format only. It cannot handle UUCP "bang paths", though simple two-component bang paths can

be converted by a straightforward rewriting configuration. This restriction does not prevent Exim from being interfaced to UUCP as a transport mechanism, provided that domain addresses are used.

- Exim insists that every address it handles has a domain attached. For incoming local messages, domainless addresses are automatically qualified with a configured domain value. Configuration options specify from which remote systems unqualified addresses are acceptable. These are then qualified on arrival.
- The only external transport mechanisms that are currently implemented are SMTP and LMTP over a TCP/IP network (including support for IPv6). However, a pipe transport is available, and there are facilities for writing messages to files and pipes, optionally in *batched SMTP* format; these facilities can be used to send messages to other transport mechanisms such as UUCP, provided they can handle domain-style addresses. Batched SMTP input is also catered for.
- Exim is not designed for storing mail for dial-in hosts. When the volumes of such mail are large, it is better to get the messages “delivered” into files (that is, off Exim’s queue) and subsequently passed on to the dial-in hosts by other means.
- Although Exim does have basic facilities for scanning incoming messages, these are not comprehensive enough to do full virus or spam scanning. Such operations are best carried out using additional specialized software packages. If you compile Exim with the content-scanning extension, straightforward interfaces to a number of common scanners are provided.

## 1.7 Run time configuration

Exim’s run time configuration is held in a single text file that is divided into a number of sections. The entries in this file consist of keywords and values, in the style of Smail 3 configuration files. A default configuration file which is suitable for simple online installations is provided in the distribution, and is described in chapter 7 below.

## 1.8 Calling interface

Like many MTAs, Exim has adopted the Sendmail command line interface so that it can be a straight replacement for */usr/lib/sendmail* or */usr/sbin/sendmail* when sending mail, but you do not need to know anything about Sendmail in order to run Exim. For actions other than sending messages, Sendmail-compatible options also exist, but those that produce output (for example, **-bp**, which lists the messages on the queue) do so in Exim’s own format. There are also some additional options that are compatible with Smail 3, and some further options that are new to Exim. Chapter 5 documents all Exim’s command line options. This information is automatically made into the man page that forms part of the Exim distribution.

Control of messages on the queue can be done via certain privileged command line options. There is also an optional monitor program called *eximon*, which displays current information in an X window, and which contains a menu interface to Exim’s command line administration options.

## 1.9 Terminology

The *body* of a message is the actual data that the sender wants to transmit. It is the last part of a message, and is separated from the *header* (see below) by a blank line.

When a message cannot be delivered, it is normally returned to the sender in a delivery failure message or a “non-delivery report” (NDR). The term *bounce* is commonly used for this action, and the error reports are often called *bounce messages*. This is a convenient shorthand for “delivery failure error report”. Such messages have an empty sender address in the message’s *envelope* (see below) to ensure that they cannot themselves give rise to further bounce messages.

The term *default* appears frequently in this manual. It is used to qualify a value which is used in the absence of any setting in the configuration. It may also qualify an action which is taken unless a configuration setting specifies otherwise.

The term *defer* is used when the delivery of a message to a specific destination cannot immediately take place for some reason (a remote host may be down, or a user's local mailbox may be full). Such deliveries are *deferred* until a later time.

The word *domain* is sometimes used to mean all but the first component of a host's name. It is *not* used in that sense here, where it normally refers to the part of an email address following the @ sign.

A message in transit has an associated *envelope*, as well as a header and a body. The envelope contains a sender address (to which bounce messages should be delivered), and any number of recipient addresses. References to the sender or the recipients of a message usually mean the addresses in the envelope. An MTA uses these addresses for delivery, and for returning bounce messages, not the addresses that appear in the header lines.

The *header* of a message is the first part of a message's text, consisting of a number of lines, each of which has a name such as *From:*, *To:*, *Subject:*, etc. Long header lines can be split over several text lines by indenting the continuations. The header is separated from the body by a blank line.

The term *local part*, which is taken from RFC 2822, is used to refer to that part of an email address that precedes the @ sign. The part that follows the @ sign is called the *domain* or *mail domain*.

The terms *local delivery* and *remote delivery* are used to distinguish delivery to a file or a pipe on the local host from delivery by SMTP over TCP/IP to another host. As far as Exim is concerned, all hosts other than the host it is running on are *remote*.

*Return path* is another name that is used for the sender address in a message's envelope.

The term *queue* is used to refer to the set of messages awaiting delivery, because this term is in widespread use in the context of MTAs. However, in Exim's case the reality is more like a pool than a queue, because there is normally no ordering of waiting messages.

The term *queue runner* is used to describe a process that scans the queue and attempts to deliver those messages whose retry times have come. This term is used by other MTAs, and also relates to the command **runq**, but in Exim the waiting messages are normally processed in an unpredictable order.

The term *spool directory* is used for a directory in which Exim keeps the messages on its queue – that is, those that it is in the process of delivering. This should not be confused with the directory in which local mailboxes are stored, which is called a “spool directory” by some people. In the Exim documentation, “spool” is always used in the first sense.

## 2. Incorporated code

A number of pieces of external code are included in the Exim distribution.

- Regular expressions are supported in the main Exim program and in the Exim monitor using the freely-distributable PCRE library, copyright © University of Cambridge. The source to PCRE is no longer shipped with Exim, so you will need to use the version of PCRE shipped with your system, or obtain and install the full version of the library from <ftp://ftp.csx.cam.ac.uk/pub/software/programming/pcre>.
- Support for the cdb (Constant DataBase) lookup method is provided by code contributed by Nigel Metheringham of (at the time he contributed it) Planet Online Ltd. The implementation is completely contained within the code of Exim. It does not link against an external cdb library. The code contains the following statements:

Copyright © 1998 Nigel Metheringham, Planet Online Ltd

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version. This code implements Dan Bernstein's Constant DataBase (cdb) spec. Information, the spec and sample code for cdb can be obtained from <http://www.pobox.com/~djb/cdb.html>. This implementation borrows some code from Dan Bernstein's implementation (which has no license restrictions applied to it).

- Client support for Microsoft's *Secure Password Authentication* is provided by code contributed by Marc Prud'hommeaux. Server support was contributed by Tom Kistner. This includes code taken from the Samba project, which is released under the Gnu GPL.
- Support for calling the Cyrus *pwcheck* and *saslauthd* daemons is provided by code taken from the Cyrus-SASL library and adapted by Alexander S. Sabourenkov. The permission notice appears below, in accordance with the conditions expressed therein.

Copyright © 2001 Carnegie Mellon University. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- (1) Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- (2) Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- (3) The name "Carnegie Mellon University" must not be used to endorse or promote products derived from this software without prior written permission. For permission or any other legal details, please contact

Office of Technology Transfer  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA 15213-3890  
(412) 268-4387, fax: (412) 268-7395  
[tech-transfer@andrew.cmu.edu](mailto:tech-transfer@andrew.cmu.edu)

- (4) Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by Computing Services at Carnegie Mellon University (<http://www.cmu.edu/computing/>)."

CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES

OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

- The Exim Monitor program, which is an X-Window application, includes modified versions of the Athena StripChart and TextPop widgets. This code is copyright by DEC and MIT, and their permission notice appears below, in accordance with the conditions expressed therein.

Copyright 1987, 1988 by Digital Equipment Corporation, Maynard, Massachusetts, and the Massachusetts Institute of Technology, Cambridge, Massachusetts.

All Rights Reserved

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the names of Digital or MIT not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

DIGITAL DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL DIGITAL BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

- The DMARC implementation uses the OpenDMARC library which is Copyrighted by The Trusted Domain Project. Portions of Exim source which use OpenDMARC derived code are indicated in the respective source files. The full OpenDMARC license is provided in the LICENSE.opendmarc file contained in the distributed source code.
- Many people have contributed code fragments, some large, some small, that were not covered by any specific licence requirements. It is assumed that the contributors are happy to see their code incorporated into Exim under the GPL.

## 3. How Exim receives and delivers mail

### 3.1 Overall philosophy

Exim is designed to work efficiently on systems that are permanently connected to the Internet and are handling a general mix of mail. In such circumstances, most messages can be delivered immediately. Consequently, Exim does not maintain independent queues of messages for specific domains or hosts, though it does try to send several messages in a single SMTP connection after a host has been down, and it also maintains per-host retry information.

### 3.2 Policy control

Policy controls are now an important feature of MTAs that are connected to the Internet. Perhaps their most important job is to stop MTAs being abused as “open relays” by misguided individuals who send out vast amounts of unsolicited junk, and want to disguise its source. Exim provides flexible facilities for specifying policy controls on incoming mail:

- Exim 4 (unlike previous versions of Exim) implements policy controls on incoming mail by means of *Access Control Lists* (ACLs). Each list is a series of statements that may either grant or deny access. ACLs can be used at several places in the SMTP dialogue while receiving a message from a remote host. However, the most common places are after each RCPT command, and at the very end of the message. The sysadmin can specify conditions for accepting or rejecting individual recipients or the entire message, respectively, at these two points (see chapter 43). Denial of access results in an SMTP error code.
- An ACL is also available for locally generated, non-SMTP messages. In this case, the only available actions are to accept or deny the entire message.
- When Exim is compiled with the content-scanning extension, facilities are provided in the ACL mechanism for passing the message to external virus and/or spam scanning software. The result of such a scan is passed back to the ACL, which can then use it to decide what to do with the message.
- When a message has been received, either from a remote host or from the local host, but before the final acknowledgment has been sent, a locally supplied C function called *local\_scan()* can be run to inspect the message and decide whether to accept it or not (see chapter 45). If the message is accepted, the list of recipients can be modified by the function.
- Using the *local\_scan()* mechanism is another way of calling external scanner software. The **SA-Exim** add-on package works this way. It does not require Exim to be compiled with the content-scanning extension.
- After a message has been accepted, a further checking mechanism is available in the form of the *system filter* (see chapter 46). This runs at the start of every delivery process.

### 3.3 User filters

In a conventional Exim configuration, users are able to run private filters by setting up appropriate *.forward* files in their home directories. See chapter 22 (about the *redirect* router) for the configuration needed to support this, and the separate document entitled *Exim's interfaces to mail filtering* for user details. Two different kinds of filtering are available:

- Sieve filters are written in the standard filtering language that is defined by RFC 3028.
- Exim filters are written in a syntax that is unique to Exim, but which is more powerful than Sieve, which it pre-dates.

User filters are run as part of the routing process, described below.

### 3.4 Message identification

Every message handled by Exim is given a *message id* which is sixteen characters long. It is divided into three parts, separated by hyphens, for example 16VDhn-0001b0-D3. Each part is a sequence of letters and digits, normally encoding numbers in base 62. However, in the Darwin operating system (Mac OS X) and when Exim is compiled to run under Cygwin, base 36 (avoiding the use of lower case letters) is used instead, because the message id is used to construct file names, and the names of files in those systems are not always case-sensitive.

The detail of the contents of the message id have changed as Exim has evolved. Earlier versions relied on the operating system not re-using a process id (pid) within one second. On modern operating systems, this assumption can no longer be made, so the algorithm had to be changed. To retain backward compatibility, the format of the message id was retained, which is why the following rules are somewhat eccentric:

- The first six characters of the message id are the time at which the message started to be received, to a granularity of one second. That is, this field contains the number of seconds since the start of the epoch (the normal Unix way of representing the date and time of day).
- After the first hyphen, the next six characters are the id of the process that received the message.
- There are two different possibilities for the final two characters:
  - (1) If **localhost\_number** is not set, this value is the fractional part of the time of reception, normally in units of 1/2000 of a second, but for systems that must use base 36 instead of base 62 (because of case-insensitive file systems), the units are 1/1000 of a second.
  - (2) If **localhost\_number** is set, it is multiplied by 200 (100) and added to the fractional part of the time, which in this case is in units of 1/200 (1/100) of a second.

After a message has been received, Exim waits for the clock to tick at the appropriate resolution before proceeding, so that if another message is received by the same process, or by another process with the same (re-used) pid, it is guaranteed that the time will be different. In most cases, the clock will already have ticked while the message was being received.

### 3.5 Receiving mail

The only way Exim can receive mail from another host is using SMTP over TCP/IP, in which case the sender and recipient addresses are transferred using SMTP commands. However, from a locally running process (such as a user's MUA), there are several possibilities:

- If the process runs Exim with the **-bm** option, the message is read non-interactively (usually via a pipe), with the recipients taken from the command line, or from the body of the message if **-t** is also used.
- If the process runs Exim with the **-bS** option, the message is also read non-interactively, but in this case the recipients are listed at the start of the message in a series of SMTP RCPT commands, terminated by a DATA command. This is so-called “batch SMTP” format, but it isn't really SMTP. The SMTP commands are just another way of passing envelope addresses in a non-interactive submission.
- If the process runs Exim with the **-bs** option, the message is read interactively, using the SMTP protocol. A two-way pipe is normally used for passing data between the local process and the Exim process. This is “real” SMTP and is handled in the same way as SMTP over TCP/IP. For example, the ACLs for SMTP commands are used for this form of submission.
- A local process may also make a TCP/IP call to the host's loopback address (127.0.0.1) or any other of its IP addresses. When receiving messages, Exim does not treat the loopback address specially. It treats all such connections in the same way as connections from other hosts.

In the three cases that do not involve TCP/IP, the sender address is constructed from the login name of the user that called Exim and a default qualification domain (which can be set by the **qualify\_domain** configuration option). For local or batch SMTP, a sender address that is passed using the SMTP MAIL command is ignored. However, the system administrator may allow certain users (“trusted

users”) to specify a different sender address unconditionally, or all users to specify certain forms of different sender address. The **-f** option or the SMTP MAIL command is used to specify these different addresses. See section 5.2 for details of trusted users, and the **untrusted\_set\_sender** option for a way of allowing untrusted users to change sender addresses.

Messages received by either of the non-interactive mechanisms are subject to checking by the non-SMTP ACL, if one is defined. Messages received using SMTP (either over TCP/IP, or interacting with a local process) can be checked by a number of ACLs that operate at different times during the SMTP session. Either individual recipients, or the entire message, can be rejected if local policy requirements are not met. The *local\_scan()* function (see chapter 45) is run for all incoming messages.

Exim can be configured not to start a delivery process when a message is received; this can be unconditional, or depend on the number of incoming SMTP connections or the system load. In these situations, new messages wait on the queue until a queue runner process picks them up. However, in standard configurations under normal conditions, delivery is started as soon as a message is received.

## 3.6 Handling an incoming message

When Exim accepts a message, it writes two files in its spool directory. The first contains the envelope information, the current status of the message, and the header lines, and the second contains the body of the message. The names of the two spool files consist of the message id, followed by **-H** for the file containing the envelope and header, and **-D** for the data file.

By default all these message files are held in a single directory called *input* inside the general Exim spool directory. Some operating systems do not perform very well if the number of files in a directory gets large; to improve performance in such cases, the **split\_spool\_directory** option can be used. This causes Exim to split up the input files into 62 sub-directories whose names are single letters or digits. When this is done, the queue is processed one sub-directory at a time instead of all at once, which can improve overall performance even when there are not enough files in each directory to affect file system performance.

The envelope information consists of the address of the message’s sender and the addresses of the recipients. This information is entirely separate from any addresses contained in the header lines. The status of the message includes a list of recipients who have already received the message. The format of the first spool file is described in chapter 56.

Address rewriting that is specified in the rewrite section of the configuration (see chapter 31) is done once and for all on incoming addresses, both in the header lines and the envelope, at the time the message is accepted. If during the course of delivery additional addresses are generated (for example, via aliasing), these new addresses are rewritten as soon as they are generated. At the time a message is actually delivered (transported) further rewriting can take place; because this is a transport option, it can be different for different forms of delivery. It is also possible to specify the addition or removal of certain header lines at the time the message is delivered (see chapters 15 and 24).

## 3.7 Life of a message

A message remains in the spool directory until it is completely delivered to its recipients or to an error address, or until it is deleted by an administrator or by the user who originally created it. In cases when delivery cannot proceed – for example, when a message can neither be delivered to its recipients nor returned to its sender, the message is marked “frozen” on the spool, and no more deliveries are attempted.

An administrator can “thaw” such messages when the problem has been corrected, and can also freeze individual messages by hand if necessary. In addition, an administrator can force a delivery error, causing a bounce message to be sent.

There are options called **ignore\_bounce\_errors\_after** and **timeout\_frozen\_after**, which discard frozen messages after a certain time. The first applies only to frozen bounces, the second to any frozen messages.

While Exim is working on a message, it writes information about each delivery attempt to its main log file. This includes successful, unsuccessful, and delayed deliveries for each recipient (see chapter



52). The log lines are also written to a separate *message log* file for each message. These logs are solely for the benefit of the administrator, and are normally deleted along with the spool files when processing of a message is complete. The use of individual message logs can be disabled by setting **no\_message\_logs**; this might give an improvement in performance on very busy systems.

All the information Exim itself needs to set up a delivery is kept in the first spool file, along with the header lines. When a successful delivery occurs, the address is immediately written at the end of a journal file, whose name is the message id followed by -J. At the end of a delivery run, if there are some addresses left to be tried again later, the first spool file (the -H file) is updated to indicate which these are, and the journal file is then deleted. Updating the spool file is done by writing a new file and renaming it, to minimize the possibility of data loss.

Should the system or the program crash after a successful delivery but before the spool file has been updated, the journal is left lying around. The next time Exim attempts to deliver the message, it reads the journal file and updates the spool file before proceeding. This minimizes the chances of double deliveries caused by crashes.

### 3.8 Processing an address for delivery

The main delivery processing elements of Exim are called *routers* and *transports*, and collectively these are known as *drivers*. Code for a number of them is provided in the source distribution, and compile-time options specify which ones are included in the binary. Run time options specify which ones are actually used for delivering messages.

Each driver that is specified in the run time configuration is an *instance* of that particular driver type. Multiple instances are allowed; for example, you can set up several different *smtp* transports, each with different option values that might specify different ports or different timeouts. Each instance has its own identifying name. In what follows we will normally use the instance name when discussing one particular instance (that is, one specific configuration of the driver), and the generic driver name when discussing the driver's features in general.

A *router* is a driver that operates on an address, either determining how its delivery should happen, by assigning it to a specific transport, or converting the address into one or more new addresses (for example, via an alias file). A router may also explicitly choose to fail an address, causing it to be bounced.

A *transport* is a driver that transmits a copy of the message from Exim's spool to some destination. There are two kinds of transport: for a *local* transport, the destination is a file or a pipe on the local host, whereas for a *remote* transport the destination is some other host. A message is passed to a specific transport as a result of successful routing. If a message has several recipients, it may be passed to a number of different transports.

An address is processed by passing it to each configured router instance in turn, subject to certain preconditions, until a router accepts the address or specifies that it should be bounced. We will describe this process in more detail shortly. First, as a simple example, we consider how each recipient address in a message is processed in a small configuration of three routers.

To make this a more concrete example, it is described in terms of some actual routers, but remember, this is only an example. You can configure Exim's routers in many different ways, and there may be any number of routers in a configuration.

The first router that is specified in a configuration is often one that handles addresses in domains that are not recognized specially by the local host. These are typically addresses for arbitrary domains on the Internet. A precondition is set up which looks for the special domains known to the host (for example, its own domain name), and the router is run for addresses that do *not* match. Typically, this is a router that looks up domains in the DNS in order to find the hosts to which this address routes. If it succeeds, the address is assigned to a suitable SMTP transport; if it does not succeed, the router is configured to fail the address.

The second router is reached only when the domain is recognized as one that "belongs" to the local host. This router does redirection – also known as aliasing and forwarding. When it generates one or more new addresses from the original, each of them is routed independently from the start. Otherwise,

the router may cause an address to fail, or it may simply decline to handle the address, in which case the address is passed to the next router.

The final router in many configurations is one that checks to see if the address belongs to a local mailbox. The precondition may involve a check to see if the local part is the name of a login account, or it may look up the local part in a file or a database. If its preconditions are not met, or if the router declines, we have reached the end of the routers. When this happens, the address is bounced.

### 3.9 Processing an address for verification

As well as being used to decide how to deliver to an address, Exim's routers are also used for *address verification*. Verification can be requested as one of the checks to be performed in an ACL for incoming messages, on both sender and recipient addresses, and it can be tested using the **-bv** and **-bvs** command line options.

When an address is being verified, the routers are run in “verify mode”. This does not affect the way the routers work, but it is a state that can be detected. By this means, a router can be skipped or made to behave differently when verifying. A common example is a configuration in which the first router sends all messages to a message-scanning program, unless they have been previously scanned. Thus, the first router accepts all addresses without any checking, making it useless for verifying. Normally, the **no\_verify** option would be set for such a router, causing it to be skipped in verify mode.

### 3.10 Running an individual router

As explained in the example above, a number of preconditions are checked before running a router. If any are not met, the router is skipped, and the address is passed to the next router. When all the preconditions on a router *are* met, the router is run. What happens next depends on the outcome, which is one of the following:

- *accept*: The router accepts the address, and either assigns it to a transport, or generates one or more “child” addresses. Processing the original address ceases, unless the **unseen** option is set on the router. This option can be used to set up multiple deliveries with different routing (for example, for keeping archive copies of messages). When **unseen** is set, the address is passed to the next router. Normally, however, an *accept* return marks the end of routing.

Any child addresses generated by the router are processed independently, starting with the first router by default. It is possible to change this by setting the **redirect\_router** option to specify which router to start at for child addresses. Unlike **pass\_router** (see below) the router specified by **redirect\_router** may be anywhere in the router configuration.

- *pass*: The router recognizes the address, but cannot handle it itself. It requests that the address be passed to another router. By default the address is passed to the next router, but this can be changed by setting the **pass\_router** option. However, (unlike **redirect\_router**) the named router must be below the current router (to avoid loops).
- *decline*: The router declines to accept the address because it does not recognize it at all. By default, the address is passed to the next router, but this can be prevented by setting the **no\_more** option. When **no\_more** is set, all the remaining routers are skipped. In effect, **no\_more** converts *decline* into *fail*.
- *fail*: The router determines that the address should fail, and queues it for the generation of a bounce message. There is no further processing of the original address unless **unseen** is set on the router.
- *defer*: The router cannot handle the address at the present time. (A database may be offline, or a DNS lookup may have timed out.) No further processing of the address happens in this delivery attempt. It is tried again next time the message is considered for delivery.
- *error*: There is some error in the router (for example, a syntax error in its configuration). The action is as for defer.

If an address reaches the end of the routers without having been accepted by any of them, it is bounced as unroutable. The default error message in this situation is “unrouteable address”, but you

can set your own message by making use of the **cannot\_route\_message** option. This can be set for any router; the value from the last router that “saw” the address is used.

Sometimes while routing you want to fail a delivery when some conditions are met but others are not, instead of passing the address on for further routing. You can do this by having a second router that explicitly fails the delivery when the relevant conditions are met. The *redirect* router has a “fail” facility for this purpose.

### 3.11 Duplicate addresses

Once routing is complete, Exim scans the addresses that are assigned to local and remote transports, and discards any duplicates that it finds. During this check, local parts are treated as case-sensitive. This happens only when actually delivering a message; when testing routers with **-bt**, all the routed addresses are shown.

### 3.12 Router preconditions

The preconditions that are tested for each router are listed below, in the order in which they are tested. The individual configuration options are described in more detail in chapter 15.

- The **local\_part\_prefix** and **local\_part\_suffix** options can specify that the local parts handled by the router may or must have certain prefixes and/or suffixes. If a mandatory affix (prefix or suffix) is not present, the router is skipped. These conditions are tested first. When an affix is present, it is removed from the local part before further processing, including the evaluation of any other conditions.
- Routers can be designated for use only when not verifying an address, that is, only when routing it for delivery (or testing its delivery routing). If the **verify** option is set false, the router is skipped when Exim is verifying an address. Setting the **verify** option actually sets two options, **verify\_sender** and **verify\_recipient**, which independently control the use of the router for sender and recipient verification. You can set these options directly if you want a router to be used for only one type of verification. Note that cutthrough delivery is classed as a recipient verification for this purpose.
- If the **address\_test** option is set false, the router is skipped when Exim is run with the **-bt** option to test an address routing. This can be helpful when the first router sends all new messages to a scanner of some sort; it makes it possible to use **-bt** to test subsequent delivery routing without having to simulate the effect of the scanner.
- Routers can be designated for use only when verifying an address, as opposed to routing it for delivery. The **verify\_only** option controls this. Again, cutthrough delivery counts as a verification.
- Individual routers can be explicitly skipped when running the routers to check an address given in the SMTP EXPN command (see the **expn** option).
- If the **domains** option is set, the domain of the address must be in the set of domains that it defines.
- If the **local\_parts** option is set, the local part of the address must be in the set of local parts that it defines. If **local\_part\_prefix** or **local\_part\_suffix** is in use, the prefix or suffix is removed from the local part before this check. If you want to do precondition tests on local parts that include affixes, you can do so by using a **condition** option (see below) that uses the variables *\$local\_part*, *\$local\_part\_prefix*, and *\$local\_part\_suffix* as necessary.
- If the **check\_local\_user** option is set, the local part must be the name of an account on the local host. If this check succeeds, the uid and gid of the local user are placed in *\$local\_user\_uid* and *\$local\_user\_gid* and the user’s home directory is placed in *\$home*; these values can be used in the remaining preconditions.
- If the **router\_home\_directory** option is set, it is expanded at this point, because it overrides the value of *\$home*. If this expansion were left till later, the value of *\$home* as set by **check\_local\_user** would be used in subsequent tests. Having two different values of *\$home* in the same router could lead to confusion.

- If the **senders** option is set, the envelope sender address must be in the set of addresses that it defines.
- If the **require\_files** option is set, the existence or non-existence of specified files is tested.
- If the **condition** option is set, it is evaluated and tested. This option uses an expanded string to allow you to set up your own custom preconditions. Expanded strings are described in chapter 11.

Note that **require\_files** comes near the end of the list, so you cannot use it to check for the existence of a file in which to lookup up a domain, local part, or sender. However, as these options are all expanded, you can use the **exists** expansion condition to make such tests within each condition. The **require\_files** option is intended for checking files that the router may be going to use internally, or which are needed by a specific transport (for example, *.procmailrc*).

### 3.13 Delivery in detail

When a message is to be delivered, the sequence of events is as follows:

- If a system-wide filter file is specified, the message is passed to it. The filter may add recipients to the message, replace the recipients, discard the message, cause a new message to be generated, or cause the message delivery to fail. The format of the system filter file is the same as for Exim user filter files, described in the separate document entitled *Exim's interfaces to mail filtering*. (**Note:** Sieve cannot be used for system filter files.)

Some additional features are available in system filters – see chapter 46 for details. Note that a message is passed to the system filter only once per delivery attempt, however many recipients it has. However, if there are several delivery attempts because one or more addresses could not be immediately delivered, the system filter is run each time. The filter condition **first\_delivery** can be used to detect the first run of the system filter.

- Each recipient address is offered to each configured router in turn, subject to its preconditions, until one is able to handle it. If no router can handle the address, that is, if they all decline, the address is failed. Because routers can be targeted at particular domains, several locally handled domains can be processed entirely independently of each other.
- A router that accepts an address may assign it to a local or a remote transport. However, the transport is not run at this time. Instead, the address is placed on a list for the particular transport, which will be run later. Alternatively, the router may generate one or more new addresses (typically from alias, forward, or filter files). New addresses are fed back into this process from the top, but in order to avoid loops, a router ignores any address which has an identically-named ancestor that was processed by itself.
- When all the routing has been done, addresses that have been successfully handled are passed to their assigned transports. When local transports are doing real local deliveries, they handle only one address at a time, but if a local transport is being used as a pseudo-remote transport (for example, to collect batched SMTP messages for transmission by some other means) multiple addresses can be handled. Remote transports can always handle more than one address at a time, but can be configured not to do so, or to restrict multiple addresses to the same domain.
- Each local delivery to a file or a pipe runs in a separate process under a non-privileged uid, and these deliveries are run one at a time. Remote deliveries also run in separate processes, normally under a uid that is private to Exim (“the Exim user”), but in this case, several remote deliveries can be run in parallel. The maximum number of simultaneous remote deliveries for any one message is set by the **remote\_max\_parallel** option. The order in which deliveries are done is not defined, except that all local deliveries happen before any remote deliveries.
- When it encounters a local delivery during a queue run, Exim checks its retry database to see if there has been a previous temporary delivery failure for the address before running the local transport. If there was a previous failure, Exim does not attempt a new delivery until the retry time for the address is reached. However, this happens only for delivery attempts that are part of a queue run. Local deliveries are always attempted when delivery immediately follows message reception, even if retry times are set for them. This makes for better behaviour if one particular message is causing problems (for example, causing quota overflow, or provoking an error in a filter file).

- Remote transports do their own retry handling, since an address may be deliverable to one of a number of hosts, each of which may have a different retry time. If there have been previous temporary failures and no host has reached its retry time, no delivery is attempted, whether in a queue run or not. See chapter 32 for details of retry strategies.
- If there were any permanent errors, a bounce message is returned to an appropriate address (the sender in the common case), with details of the error for each failing address. Exim can be configured to send copies of bounce messages to other addresses.
- If one or more addresses suffered a temporary failure, the message is left on the queue, to be tried again later. Delivery of these addresses is said to be *deferred*.
- When all the recipient addresses have either been delivered or bounced, handling of the message is complete. The spool files and message log are deleted, though the message log can optionally be preserved if required.

### 3.14 Retry mechanism

Exim's mechanism for retrying messages that fail to get delivered at the first attempt is the queue runner process. You must either run an Exim daemon that uses the **-q** option with a time interval to start queue runners at regular intervals, or use some other means (such as *cron*) to start them. If you do not arrange for queue runners to be run, messages that fail temporarily at the first attempt will remain on your queue for ever. A queue runner process works its way through the queue, one message at a time, trying each delivery that has passed its retry time. You can run several queue runners at once.

Exim uses a set of configured rules to determine when next to retry the failing address (see chapter 32). These rules also specify when Exim should give up trying to deliver to the address, at which point it generates a bounce message. If no retry rules are set for a particular host, address, and error combination, no retries are attempted, and temporary errors are treated as permanent.

### 3.15 Temporary delivery failure

There are many reasons why a message may not be immediately deliverable to a particular address. Failure to connect to a remote machine (because it, or the connection to it, is down) is one of the most common. Temporary failures may be detected during routing as well as during the transport stage of delivery. Local deliveries may be delayed if NFS files are unavailable, or if a mailbox is on a file system where the user is over quota. Exim can be configured to impose its own quotas on local mailboxes; where system quotas are set they will also apply.

If a host is unreachable for a period of time, a number of messages may be waiting for it by the time it recovers, and sending them in a single SMTP connection is clearly beneficial. Whenever a delivery to a remote host is deferred, Exim makes a note in its hints database, and whenever a successful SMTP delivery has happened, it looks to see if any other messages are waiting for the same host. If any are found, they are sent over the same SMTP connection, subject to a configuration limit as to the maximum number in any one connection.

### 3.16 Permanent delivery failure

When a message cannot be delivered to some or all of its intended recipients, a bounce message is generated. Temporary delivery failures turn into permanent errors when their timeout expires. All the addresses that fail in a given delivery attempt are listed in a single message. If the original message has many recipients, it is possible for some addresses to fail in one delivery attempt and others to fail subsequently, giving rise to more than one bounce message. The wording of bounce messages can be customized by the administrator. See chapter 49 for details.

Bounce messages contain an *X-Failed-Recipients:* header line that lists the failed addresses, for the benefit of programs that try to analyse such messages automatically.

A bounce message is normally sent to the sender of the original message, as obtained from the message's envelope. For incoming SMTP messages, this is the address given in the MAIL command.

However, when an address is expanded via a forward or alias file, an alternative address can be specified for delivery failures of the generated addresses. For a mailing list expansion (see section 50.2) it is common to direct bounce messages to the manager of the list.

### 3.17 Failures to deliver bounce messages

If a bounce message (either locally generated or received from a remote host) itself suffers a permanent delivery failure, the message is left on the queue, but it is frozen, awaiting the attention of an administrator. There are options that can be used to make Exim discard such failed messages, or to keep them for only a short time (see **timeout\_frozen\_after** and **ignore\_bounce\_errors\_after**).

## 4. Building and installing Exim

### 4.1 Unpacking

Exim is distributed as a gzipped or bziped tar file which, when unpacked, creates a directory with the name of the current release (for example, *exim-4.90*) into which the following files are placed:

<i>ACKNOWLEDGMENTS</i>	contains some acknowledgments
<i>CHANGES</i>	contains a reference to where changes are documented
<i>LICENCE</i>	the GNU General Public Licence
<i>Makefile</i>	top-level make file
<i>NOTICE</i>	conditions for the use of Exim
<i>README</i>	list of files, directories and simple build instructions

Other files whose names begin with *README* may also be present. The following subdirectories are created:

<i>Local</i>	an empty directory for local configuration files
<i>OS</i>	OS-specific files
<i>doc</i>	documentation files
<i>exim_monitor</i>	source files for the Exim monitor
<i>scripts</i>	scripts used in the build process
<i>src</i>	remaining source files
<i>util</i>	independent utilities

The main utility programs are contained in the *src* directory, and are built with the Exim binary. The *util* directory contains a few optional scripts that may be useful to some sites.

### 4.2 Multiple machine architectures and operating systems

The building process for Exim is arranged to make it easy to build binaries for a number of different architectures and operating systems from the same set of source files. Compilation does not take place in the *src* directory. Instead, a *build directory* is created for each architecture and operating system. Symbolic links to the sources are installed in this directory, which is where the actual building takes place. In most cases, Exim can discover the machine architecture and operating system for itself, but the defaults can be overridden if necessary. A C99-capable compiler will be required for the build.

### 4.3 PCRE library

Exim no longer has an embedded PCRE library as the vast majority of modern systems include PCRE as a system library, although you may need to install the PCRE or PCRE development package for your operating system. If your system has a normal PCRE installation the Exim build process will need no further configuration. If the library or the headers are in an unusual location you will need to either set the PCRE\_LIBS and INCLUDE directives appropriately, or set PCRE\_CONFIG=yes to use the installed *pcre-config* command. If your operating system has no PCRE support then you will need to obtain and build the current PCRE from <ftp://ftp.csx.cam.ac.uk/pub/software/programming/pcre/>. More information on PCRE is available at <http://www.pcre.org/>.

### 4.4 DBM libraries

Even if you do not use any DBM files in your configuration, Exim still needs a DBM library in order to operate, because it uses indexed files for its hints databases. Unfortunately, there are a number of DBM libraries in existence, and different operating systems often have different ones installed.

If you are using Solaris, IRIX, one of the modern BSD systems, or a modern Linux distribution, the DBM configuration should happen automatically, and you may be able to ignore this section. Otherwise, you may have to learn more than you would like about DBM libraries from what follows.

Licensed versions of Unix normally contain a library of DBM functions operating via the *ndbm* interface, and this is what Exim expects by default. Free versions of Unix seem to vary in what they

contain as standard. In particular, some early versions of Linux have no default DBM library, and different distributors have chosen to bundle different libraries with their packaged versions. However, the more recent releases seem to have standardized on the Berkeley DB library.

Different DBM libraries have different conventions for naming the files they use. When a program opens a file called *dbmfile*, there are several possibilities:

- (1) A traditional *ndbm* implementation, such as that supplied as part of Solaris, operates on two files called *dbmfile.dir* and *dbmfile.pag*.
- (2) The GNU library, *gdbm*, operates on a single file. If used via its *ndbm* compatibility interface it makes two different hard links to it with names *dbmfile.dir* and *dbmfile.pag*, but if used via its native interface, the file name is used unmodified.
- (3) The Berkeley DB package, if called via its *ndbm* compatibility interface, operates on a single file called *dbmfile.db*, but otherwise looks to the programmer exactly the same as the traditional *ndbm* implementation.
- (4) If the Berkeley package is used in its native mode, it operates on a single file called *dbmfile*; the programmer's interface is somewhat different to the traditional *ndbm* interface.
- (5) To complicate things further, there are several very different versions of the Berkeley DB package. Version 1.85 was stable for a very long time, releases 2.x and 3.x were current for a while, but the latest versions are now numbered 4.x. Maintenance of some of the earlier releases has ceased. All versions of Berkeley DB can be obtained from <http://www.sleepycat.com/>.
- (6) Yet another DBM library, called *tdb*, is available from <http://download.sourceforge.net/tdb>. It has its own interface, and also operates on a single file.

Exim and its utilities can be compiled to use any of these interfaces. In order to use any version of the Berkeley DB package in native mode, you must set `USE_DB` in an appropriate configuration file (typically *Local/Makefile*). For example:

```
USE_DB=yes
```

Similarly, for *gdbm* you set `USE_GDBM`, and for *tdb* you set `USE_TDB`. An error is diagnosed if you set more than one of these.

At the lowest level, the build-time configuration sets none of these options, thereby assuming an interface of type (1). However, some operating system configuration files (for example, those for the BSD operating systems and Linux) assume type (4) by setting `USE_DB` as their default, and the configuration files for Cygwin set `USE_GDBM`. Anything you set in *Local/Makefile*, however, overrides these system defaults.

As well as setting `USE_DB`, `USE_GDBM`, or `USE_TDB`, it may also be necessary to set `DBMLIB`, to cause inclusion of the appropriate library, as in one of these lines:

```
DBMLIB = -ldb
DBMLIB = -ltdb
```

Settings like that will work if the DBM library is installed in the standard place. Sometimes it is not, and the library's header file may also not be in the default path. You may need to set `INCLUDE` to specify where the header file is, and to specify the path to the library more fully in `DBMLIB`, as in this example:

```
INCLUDE=-I/usr/local/include/db-4.1
DBMLIB=/usr/local/lib/db-4.1/libdb.a
```

There is further detailed discussion about the various DBM libraries in the file *doc/dbm.discuss.txt* in the Exim distribution.

## 4.5 Pre-building configuration

Before building Exim, a local configuration file that specifies options independent of any operating system has to be created with the name *Local/Makefile*. A template for this file is supplied as the file *src/EDITME*, and it contains full descriptions of all the option settings therein. These descriptions are



therefore not repeated here. If you are building Exim for the first time, the simplest thing to do is to copy *src/EDITME* to *Local/Makefile*, then read it and edit it appropriately.

There are three settings that you must supply, because Exim will not build without them. They are the location of the run time configuration file (`CONFIGURE_FILE`), the directory in which Exim binaries will be installed (`BIN_DIRECTORY`), and the identity of the Exim user (`EXIM_USER` and maybe `EXIM_GROUP` as well). The value of `CONFIGURE_FILE` can in fact be a colon-separated list of file names; Exim uses the first of them that exists.

There are a few other parameters that can be specified either at build time or at run time, to enable the same binary to be used on a number of different machines. However, if the locations of Exim's spool directory and log file directory (if not within the spool directory) are fixed, it is recommended that you specify them in *Local/Makefile* instead of at run time, so that errors detected early in Exim's execution (such as a malformed configuration file) can be logged.

Exim's interfaces for calling virus and spam scanning software directly from access control lists are not compiled by default. If you want to include these facilities, you need to set

```
WITH_CONTENT_SCAN=yes
```

in your *Local/Makefile*. For details of the facilities themselves, see chapter 44.

If you are going to build the Exim monitor, a similar configuration process is required. The file *exim\_monitor/EDITME* must be edited appropriately for your installation and saved under the name *Local/eximon.conf*. If you are happy with the default settings described in *exim\_monitor/EDITME*, *Local/eximon.conf* can be empty, but it must exist.

This is all the configuration that is needed in straightforward cases for known operating systems. However, the building process is set up so that it is easy to override options that are set by default or by operating-system-specific configuration files, for example to change the name of the C compiler, which defaults to **gcc**. See section 4.13 below for details of how to do this.

## 4.6 Support for iconv()

The contents of header lines in messages may be encoded according to the rules described RFC 2047. This makes it possible to transmit characters that are not in the ASCII character set, and to label them as being in a particular character set. When Exim is inspecting header lines by means of the `$h_` mechanism, it decodes them, and translates them into a specified character set (default is set at build time). The translation is possible only if the operating system supports the *iconv()* function.

However, some of the operating systems that supply *iconv()* do not support very many conversions. The GNU **libiconv** library (available from <http://www.gnu.org/software/libiconv/>) can be installed on such systems to remedy this deficiency, as well as on systems that do not supply *iconv()* at all. After installing **libiconv**, you should add

```
HAVE_ICONV=yes
```

to your *Local/Makefile* and rebuild Exim.

## 4.7 Including TLS/SSL encryption support

Exim can be built to support encrypted SMTP connections, using the STARTTLS command as per RFC 2487. It can also support legacy clients that expect to start a TLS session immediately on connection to a non-standard port (see the `tls_on_connect_ports` runtime option and the `-tls-on-connect` command line option).

If you want to build Exim with TLS support, you must first install either the OpenSSL or GnuTLS library. There is no cryptographic code in Exim itself for implementing SSL.

If OpenSSL is installed, you should set

```
SUPPORT_TLS=yes
TLS_LIBS=-lssl -lcrypto
```

in *Local/Makefile*. You may also need to specify the locations of the OpenSSL library and include files. For example:

```
SUPPORT_TLS=yes
TLS_LIBS=-L/usr/local/openssl/lib -lssl -lcrypto
TLS_INCLUDE=-I/usr/local/openssl/include/
```

If you have *pkg-config* available, then instead you can just use:

```
SUPPORT_TLS=yes
USE_OPENSSL_PC=openssl
```

If GnuTLS is installed, you should set

```
SUPPORT_TLS=yes
USE_GNUTLS=yes
TLS_LIBS=-lgnutls -ltasn1 -lgcrypt
```

in *Local/Makefile*, and again you may need to specify the locations of the library and include files. For example:

```
SUPPORT_TLS=yes
USE_GNUTLS=yes
TLS_LIBS=-L/usr/gnu/lib -lgnutls -ltasn1 -lgcrypt
TLS_INCLUDE=-I/usr/gnu/include
```

If you have *pkg-config* available, then instead you can just use:

```
SUPPORT_TLS=yes
USE_GNUTLS=yes
USE_GNUTLS_PC=gnutls
```

You do not need to set `TLS_INCLUDE` if the relevant directory is already specified in `INCLUDE`. Details of how to configure Exim to make use of TLS are given in chapter 42.

## 4.8 Use of *tcpwrappers*

Exim can be linked with the *tcpwrappers* library in order to check incoming SMTP calls using the *tcpwrappers* control files. This may be a convenient alternative to Exim's own checking facilities for installations that are already making use of *tcpwrappers* for other purposes. To do this, you should set `USE_TCP_WRAPPERS` in *Local/Makefile*, arrange for the file *tcpd.h* to be available at compile time, and also ensure that the library *libwrap.a* is available at link time, typically by including **-lwrap** in `EXTRALIBS_EXIM`. For example, if *tcpwrappers* is installed in */usr/local*, you might have

```
USE_TCP_WRAPPERS=yes
CFLAGS=-O -I/usr/local/include
EXTRALIBS_EXIM=-L/usr/local/lib -lwrap
```

in *Local/Makefile*. The daemon name to use in the *tcpwrappers* control files is “exim”. For example, the line

```
exim : LOCAL 192.168.1. .friendly.domain.example
```

in your */etc/hosts.allow* file allows connections from the local host, from the subnet 192.168.1.0/24, and from all hosts in *friendly.domain.example*. All other connections are denied. The daemon name used by *tcpwrappers* can be changed at build time by setting `TCP_WRAPPERS_DAEMON_NAME` in *Local/Makefile*, or by setting `tcp_wrappers_daemon_name` in the configure file. Consult the *tcpwrappers* documentation for further details.

## 4.9 Including support for IPv6

Exim contains code for use on systems that have IPv6 support. Setting `HAVE_IPV6=YES` in *Local/Makefile* causes the IPv6 code to be included; it may also be necessary to set `IPV6_INCLUDE` and `IPV6_LIBS` on systems where the IPv6 support is not fully integrated into the normal include and library files.

Two different types of DNS record for handling IPv6 addresses have been defined. AAAA records (analogous to A records for IPv4) are in use, and are currently seen as the mainstream. Another record type called A6 was proposed as better than AAAA because it had more flexibility. However, it was felt to be over-complex, and its status was reduced to “experimental”. Exim used to have a compile option for including A6 record support but this has now been withdrawn.

## 4.10 Dynamically loaded lookup module support

On some platforms, Exim supports not compiling all lookup types directly into the main binary, instead putting some into external modules which can be loaded on demand. This permits packagers to build Exim with support for lookups with extensive library dependencies without requiring all users to install all of those dependencies. Most, but not all, lookup types can be built this way.

Set `LOOKUP_MODULE_DIR` to the directory into which the modules will be installed; Exim will only load modules from that directory, as a security measure. You will need to set `CFLAGS_DYNAMIC` if not already defined for your OS; see *OS/Makefile-Linux* for an example. Some other requirements for adjusting `EXTRALIBS` may also be necessary, see *src/EDITME* for details.

Then, for each module to be loaded dynamically, define the relevant `LOOKUP_<lookup_type>` flags to have the value “2” instead of “yes”. For example, this will build in `lsearch` but load `sqlite` and `mysql` support on demand:

```
LOOKUP_LSEARCH=yes
LOOKUP_SQLITE=2
LOOKUP_MYSQL=2
```

## 4.11 The building process

Once *Local/Makefile* (and *Local/eximon.conf*, if required) have been created, run *make* at the top level. It determines the architecture and operating system types, and creates a build directory if one does not exist. For example, on a Sun system running Solaris 8, the directory *build-SunOS5-5.8-sparc* is created. Symbolic links to relevant source files are installed in the build directory.

If this is the first time *make* has been run, it calls a script that builds a make file inside the build directory, using the configuration files from the *Local* directory. The new make file is then passed to another instance of *make*. This does the real work, building a number of utility scripts, and then compiling and linking the binaries for the Exim monitor (if configured), a number of utility programs, and finally Exim itself. The command `make makefile` can be used to force a rebuild of the make file in the build directory, should this ever be necessary.

If you have problems building Exim, check for any comments there may be in the *README* file concerning your operating system, and also take a look at the FAQ, where some common problems are covered.

## 4.12 Output from “make”

The output produced by the *make* process for compile lines is often very unreadable, because these lines can be very long. For this reason, the normal output is suppressed by default, and instead output similar to that which appears when compiling the 2.6 Linux kernel is generated: just a short line for each module that is being compiled or linked. However, it is still possible to get the full output, by calling *make* like this:

```
FULLECHO=' ' make -e
```

The value of `FULLECHO` defaults to “@”, the flag character that suppresses command reflection in *make*. When you ask for the full output, it is given in addition to the short output.

## 4.13 Overriding build-time options for Exim

The main make file that is created at the beginning of the building process consists of the concatenation of a number of files which set configuration values, followed by a fixed set of *make* instruc-

tions. If a value is set more than once, the last setting overrides any previous ones. This provides a convenient way of overriding defaults. The files that are concatenated are, in order:

```
OS/Makefile-Default
OS/Makefile-<ostype>
Local/Makefile
Local/Makefile-<ostype>
Local/Makefile-<archtype>
Local/Makefile-<ostype>-<archtype>
OS/Makefile-Base
```

where *<ostype>* is the operating system type and *<archtype>* is the architecture type. *Local/Makefile* is required to exist, and the building process fails if it is absent. The other three *Local* files are optional, and are often not needed.

The values used for *<ostype>* and *<archtype>* are obtained from scripts called *scripts/os-type* and *scripts/arch-type* respectively. If either of the environment variables EXIM\_OSTYPE or EXIM\_ARCHTYPE is set, their values are used, thereby providing a means of forcing particular settings. Otherwise, the scripts try to get values from the **uname** command. If this fails, the shell variables OSTYPE and ARCHTYPE are inspected. A number of *ad hoc* transformations are then applied, to produce the standard names that Exim expects. You can run these scripts directly from the shell in order to find out what values are being used on your system.

*OS/Makefile-Default* contains comments about the variables that are set therein. Some (but not all) are mentioned below. If there is something that needs changing, review the contents of this file and the contents of the make file for your operating system (*OS/Makefile-<ostype>*) to see what the default values are.

If you need to change any of the values that are set in *OS/Makefile-Default* or in *OS/Makefile-<ostype>*, or to add any new definitions, you do not need to change the original files. Instead, you should make the changes by putting the new values in an appropriate *Local* file. For example, when building Exim in many releases of the Tru64-Unix (formerly Digital UNIX, formerly DEC-OSF1) operating system, it is necessary to specify that the C compiler is called *cc* rather than *gcc*. Also, the compiler must be called with the option **-std1**, to make it recognize some of the features of Standard C that Exim uses. (Most other compilers recognize Standard C by default.) To do this, you should create a file called *Local/Makefile-OSF1* containing the lines

```
CC=cc
CFLAGS=-std1
```

If you are compiling for just one operating system, it may be easier to put these lines directly into *Local/Makefile*.

Keeping all your local configuration settings separate from the distributed files makes it easy to transfer them to new versions of Exim simply by copying the contents of the *Local* directory.

Exim contains support for doing LDAP, NIS, NIS+, and other kinds of file lookup, but not all systems have these components installed, so the default is not to include the relevant code in the binary. All the different kinds of file and database lookup that Exim supports are implemented as separate code modules which are included only if the relevant compile-time options are set. In the case of LDAP, NIS, and NIS+, the settings for *Local/Makefile* are:

```
LOOKUP_LDAP=yes
LOOKUP_NIS=yes
LOOKUP_NISPLUS=yes
```

and similar settings apply to the other lookup types. They are all listed in *src/EDITME*. In many cases the relevant include files and interface libraries need to be installed before compiling Exim. However, there are some optional lookup types (such as *cdb*) for which the code is entirely contained within Exim, and no external include files or libraries are required. When a lookup type is not included in the binary, attempts to configure Exim to use it cause run time configuration errors.

Many systems now use a tool called *pkg-config* to encapsulate information about how to compile against a library; Exim has some initial support for being able to use *pkg-config* for lookups and

authenticators. For any given makefile variable which starts LOOKUP\_ or AUTH\_, you can add a new variable with the \_PC suffix in the name and assign as the value the name of the package to be queried. The results of querying via the *pkg-config* command will be added to the appropriate Makefile variables with += directives, so your version of *make* will need to support that syntax. For instance:

```
LOOKUP_SQLITE=yes
LOOKUP_SQLITE_PC=sqlite3
AUTH_GSASL=yes
AUTH_GSASL_PC=libgsasl
AUTH_HEIMDAL_GSSAPI=yes
AUTH_HEIMDAL_GSSAPI_PC=heimdal-gssapi
```

Exim can be linked with an embedded Perl interpreter, allowing Perl subroutines to be called during string expansion. To enable this facility,

```
EXIM_PERL=perl.o
```

must be defined in *Local/Makefile*. Details of this facility are given in chapter 12.

The location of the X11 libraries is something that varies a lot between operating systems, and there may be different versions of X11 to cope with. Exim itself makes no use of X11, but if you are compiling the Exim monitor, the X11 libraries must be available. The following three variables are set in *OS/Makefile-Default*:

```
X11=/usr/X11R6
XINCLUDE=-I$(X11)/include
XLFLAGS=-L$(X11)/lib
```

These are overridden in some of the operating-system configuration files. For example, in *OS/Makefile-SunOS5* there is

```
X11=/usr/openwin
XINCLUDE=-I$(X11)/include
XLFLAGS=-L$(X11)/lib -R$(X11)/lib
```

If you need to override the default setting for your operating system, place a definition of all three of these variables into your *Local/Makefile-*<ostype>** file.

If you need to add any extra libraries to the link steps, these can be put in a variable called EXTRALIBS, which appears in all the link commands, but by default is not defined. In contrast, EXTRALIBS\_EXIM is used only on the command for linking the main Exim binary, and not for any associated utilities.

There is also DBMLIB, which appears in the link commands for binaries that use DBM functions (see also section 4.4). Finally, there is EXTRALIBS\_EXIMON, which appears only in the link step for the Exim monitor binary, and which can be used, for example, to include additional X11 libraries.

The make file copes with rebuilding Exim correctly if any of the configuration files are edited. However, if an optional configuration file is deleted, it is necessary to touch the associated non-optional file (that is, *Local/Makefile* or *Local/eximon.conf*) before rebuilding.

## 4.14 OS-specific header files

The *OS* directory contains a number of files with names of the form *os.h-*<ostype>**. These are system-specific C header files that should not normally need to be changed. There is a list of macro settings that are recognized in the file *OS/os.configuring*, which should be consulted if you are porting Exim to a new operating system.

## 4.15 Overriding build-time options for the monitor

A similar process is used for overriding things when building the Exim monitor, where the files that are involved are

```
OS/eximon.conf-Default
OS/eximon.conf-<ostype>
Local/eximon.conf
Local/eximon.conf-<ostype>
Local/eximon.conf-<archtype>
Local/eximon.conf-<ostype>-<archtype>
```

As with Exim itself, the final three files need not exist, and in this case the *OS/eximon.conf-<ostype>* file is also optional. The default values in *OS/eximon.conf-Default* can be overridden dynamically by setting environment variables of the same name, preceded by *EXIMON\_*. For example, setting *EXIMON\_LOG\_DEPTH* in the environment overrides the value of *LOG\_DEPTH* at run time.

## 4.16 Installing Exim binaries and scripts

The command `make install` runs the *exim\_install* script with no arguments. The script copies binaries and utility scripts into the directory whose name is specified by the *BIN\_DIRECTORY* setting in *Local/Makefile*. The install script copies files only if they are newer than the files they are going to replace. The Exim binary is required to be owned by root and have the *setuid* bit set, for normal configurations. Therefore, you must run `make install` as root so that it can set up the Exim binary in this way. However, in some special situations (for example, if a host is doing no local deliveries) it may be possible to run Exim without making the binary *setuid* root (see chapter 55 for details).

Exim's run time configuration file is named by the *CONFIGURE\_FILE* setting in *Local/Makefile*. If this names a single file, and the file does not exist, the default configuration file *src/configure.default* is copied there by the installation script. If a run time configuration file already exists, it is left alone. If *CONFIGURE\_FILE* is a colon-separated list, naming several alternative files, no default is installed.

One change is made to the default configuration file when it is installed: the default configuration contains a router that references a system aliases file. The path to this file is set to the value specified by *SYSTEM\_ALIASES\_FILE* in *Local/Makefile* (*/etc/aliases* by default). If the system aliases file does not exist, the installation script creates it, and outputs a comment to the user.

The created file contains no aliases, but it does contain comments about the aliases a site should normally have. Mail aliases have traditionally been kept in */etc/aliases*. However, some operating systems are now using */etc/mail/aliases*. You should check if yours is one of these, and change Exim's configuration if necessary.

The default configuration uses the local host's name as the only local domain, and is set up to do local deliveries into the shared directory */var/mail*, running as the local user. System aliases and *forward* files in users' home directories are supported, but no NIS or NIS+ support is configured. Domains other than the name of the local host are routed using the DNS, with delivery over SMTP.

It is possible to install Exim for special purposes (such as building a binary distribution) in a private part of the file system. You can do this by a command such as

```
make DESTDIR=/some/directory/ install
```

This has the effect of pre-pending the specified directory to all the file paths, except the name of the system aliases file that appears in the default configuration. (If a default alias file is created, its name is modified.) For backwards compatibility, *ROOT* is used if *DESTDIR* is not set, but this usage is deprecated.

Running `make install` does not copy the Exim 4 conversion script *convert4r4*. You will probably run this only once if you are upgrading from Exim 3. None of the documentation files in the *doc* directory are copied, except for the info files when you have set *INFO\_DIRECTORY*, as described in section 4.17 below.

For the utility programs, old versions are renamed by adding the suffix *.O* to their names. The Exim binary itself, however, is handled differently. It is installed under a name that includes the version number and the compile number, for example *exim-4.90-1*. The script then arranges for a symbolic

link called *exim* to point to the binary. If you are updating a previous version of Exim, the script takes care to ensure that the name *exim* is never absent from the directory (as seen by other processes).

If you want to see what the *make install* will do before running it for real, you can pass the **-n** option to the installation script by this command:

```
make INSTALL_ARG=-n install
```

The contents of the variable `INSTALL_ARG` are passed to the installation script. You do not need to be root to run this test. Alternatively, you can run the installation script directly, but this must be from within the build directory. For example, from the top-level Exim directory you could use this command:

```
(cd build-SunOS5-5.5.1-sparc; ../scripts/exim_install -n)
```

There are two other options that can be supplied to the installation script.

- **-no\_chown** bypasses the call to change the owner of the installed binary to root, and the call to make it a setuid binary.
- **-no\_symlink** bypasses the setting up of the symbolic link *exim* to the installed binary.

`INSTALL_ARG` can be used to pass these options to the script. For example:

```
make INSTALL_ARG=-no_symlink install
```

The installation script can also be given arguments specifying which files are to be copied. For example, to install just the Exim binary, and nothing else, without creating the symbolic link, you could use:

```
make INSTALL_ARG='-no_symlink exim' install
```

## 4.17 Installing info documentation

Not all systems use the GNU *info* system for documentation, and for this reason, the Texinfo source of Exim's documentation is not included in the main distribution. Instead it is available separately from the ftp site (see section 1.5).

If you have defined `INFO_DIRECTORY` in *Local/Makefile* and the Texinfo source of the documentation is found in the source tree, running `make install` automatically builds the info files and installs them.

## 4.18 Setting up the spool directory

When it starts up, Exim tries to create its spool directory if it does not exist. The Exim uid and gid are used for the owner and group of the spool directory. Sub-directories are automatically created in the spool directory as necessary.

## 4.19 Testing

Having installed Exim, you can check that the run time configuration file is syntactically valid by running the following command, which assumes that the Exim binary directory is within your `PATH` environment variable:

```
exim -bV
```

If there are any errors in the configuration file, Exim outputs error messages. Otherwise it outputs the version number and build date, the DBM library that is being used, and information about which drivers and other optional code modules are included in the binary. Some simple routing tests can be done by using the address testing option. For example,

```
exim -bt <local username>
```

should verify that it recognizes a local mailbox, and

```
exim -bt <remote address>
```

a remote one. Then try getting it to deliver mail, both locally and remotely. This can be done by passing messages directly to Exim, without going through a user agent. For example:

```
exim -v postmaster@your.domain.example
From: user@your.domain.example
To: postmaster@your.domain.example
Subject: Testing Exim
```

```
This is a test message.
^D
```

The **-v** option causes Exim to output some verification of what it is doing. In this case you should see copies of three log lines, one for the message's arrival, one for its delivery, and one containing "Completed".

If you encounter problems, look at Exim's log files (*mainlog* and *paniclog*) to see if there is any relevant information there. Another source of information is running Exim with debugging turned on, by specifying the **-d** option. If a message is stuck on Exim's spool, you can force a delivery with debugging turned on by a command of the form

```
exim -d -M<exim-message-id>
```

You must be root or an "admin user" in order to do this. The **-d** option produces rather a lot of output, but you can cut this down to specific areas. For example, if you use **-d-all+route** only the debugging information relevant to routing is included. (See the **-d** option in chapter 5 for more details.)

One specific problem that has shown up on some sites is the inability to do local deliveries into a shared mailbox directory, because it does not have the "sticky bit" set on it. By default, Exim tries to create a lock file before writing to a mailbox file, and if it cannot create the lock file, the delivery is deferred. You can get round this either by setting the "sticky bit" on the directory, or by setting a specific group for local deliveries and allowing that group to create files in the directory (see the comments above the *local\_delivery* transport in the default configuration file). Another approach is to configure Exim not to use lock files, but just to rely on *fcntl()* locking instead. However, you should do this only if all user agents also use *fcntl()* locking. For further discussion of locking issues, see chapter 26.

One thing that cannot be tested on a system that is already running an MTA is the receipt of incoming SMTP mail on the standard SMTP port. However, the **-oX** option can be used to run an Exim daemon that listens on some other port, or *inetd* can be used to do this. The **-bh** option and the *exim\_checkaccess* utility can be used to check out policy controls on incoming SMTP mail.

Testing a new version on a system that is already running Exim can most easily be done by building a binary with a different *CONFIGURE\_FILE* setting. From within the run time configuration, all other file and directory names that Exim uses can be altered, in order to keep it entirely clear of the production version.

## 4.20 Replacing another MTA with Exim

Building and installing Exim for the first time does not of itself put it in general use. The name by which the system's MTA is called by mail user agents is either */usr/sbin/sendmail*, or */usr/lib/sendmail* (depending on the operating system), and it is necessary to make this name point to the *exim* binary in order to get the user agents to pass messages to Exim. This is normally done by renaming any existing file and making */usr/sbin/sendmail* or */usr/lib/sendmail* a symbolic link to the *exim* binary. It is a good idea to remove any setuid privilege and executable status from the old MTA. It is then necessary to stop and restart the mailer daemon, if one is running.

Some operating systems have introduced alternative ways of switching MTAs. For example, if you are running FreeBSD, you need to edit the file */etc/mail/mailer.conf* instead of setting up a symbolic link as just described. A typical example of the contents of this file for running Exim is as follows:

```
sendmail          /usr/exim/bin/exim
send-mail         /usr/exim/bin/exim
```



```
mailq          /usr/exim/bin/exim -bp
newaliases     /usr/bin/true
```

Once you have set up the symbolic link, or edited */etc/mail/mailer.conf*, your Exim installation is “live”. Check it by sending a message from your favourite user agent.

You should consider what to tell your users about the change of MTA. Exim may have different capabilities to what was previously running, and there are various operational differences such as the text of messages produced by command line options and in bounce messages. If you allow your users to make use of Exim’s filtering capabilities, you should make the document entitled *Exim’s interface to mail filtering* available to them.

## 4.21 Upgrading Exim

If you are already running Exim on your host, building and installing a new version automatically makes it available to MUAs, or any other programs that call the MTA directly. However, if you are running an Exim daemon, you do need to send it a HUP signal, to make it re-execute itself, and thereby pick up the new binary. You do not need to stop processing mail in order to install a new version of Exim. The install script does not modify an existing runtime configuration file.

## 4.22 Stopping the Exim daemon on Solaris

The standard command for stopping the mailer daemon on Solaris is

```
/etc/init.d/sendmail stop
```

If */usr/lib/sendmail* has been turned into a symbolic link, this script fails to stop Exim because it uses the command *ps -e* and greps the output for the text “sendmail”; this is not present because the actual program name (that is, “exim”) is given by the *ps* command with these options. A solution is to replace the line that finds the process id with something like

```
pid=`cat /var/spool/exim/exim-daemon.pid`
```

to obtain the daemon’s pid directly from the file that Exim saves it in.

Note, however, that stopping the daemon does not “stop Exim”. Messages can still be received from local processes, and if automatic delivery is configured (the normal case), deliveries will still occur.

## 5. The Exim command line

Exim's command line takes the standard Unix form of a sequence of options, each starting with a hyphen character, followed by a number of arguments. The options are compatible with the main options of Sendmail, and there are also some additional options, some of which are compatible with Smail 3. Certain combinations of options do not make sense, and provoke an error if used. The form of the arguments depends on which options are set.

### 5.1 Setting options by program name

If Exim is called under the name *mailq*, it behaves as if the option **-bp** were present before any other options. The **-bp** option requests a listing of the contents of the mail queue on the standard output. This feature is for compatibility with some systems that contain a command of that name in one of the standard libraries, symbolically linked to */usr/sbin/sendmail* or */usr/lib/sendmail*.

If Exim is called under the name *rsmtplib* it behaves as if the option **-bS** were present before any other options, for compatibility with Smail. The **-bS** option is used for reading in a number of messages in batched SMTP format.

If Exim is called under the name *rmail* it behaves as if the **-i** and **-oe** options were present before any other options, for compatibility with Smail. The name *rmail* is used as an interface by some UUCP systems.

If Exim is called under the name *runq* it behaves as if the option **-q** were present before any other options, for compatibility with Smail. The **-q** option causes a single queue runner process to be started.

If Exim is called under the name *newaliases* it behaves as if the option **-bi** were present before any other options, for compatibility with Sendmail. This option is used for rebuilding Sendmail's alias file. Exim does not have the concept of a single alias file, but can be configured to run a given command if called with the **-bi** option.

### 5.2 Trusted and admin users

Some Exim options are available only to *trusted users* and others are available only to *admin users*. In the description below, the phrases "Exim user" and "Exim group" mean the user and group defined by EXIM\_USER and EXIM\_GROUP in *Local/Makefile* or set by the **exim\_user** and **exim\_group** options. These do not necessarily have to use the name "exim".

- The trusted users are root, the Exim user, any user listed in the **trusted\_users** configuration option, and any user whose current group or any supplementary group is one of those listed in the **trusted\_groups** configuration option. Note that the Exim group is not automatically trusted.

Trusted users are always permitted to use the **-f** option or a leading "From " line to specify the envelope sender of a message that is passed to Exim through the local interface (see the **-bm** and **-f** options below). See the **untrusted\_set\_sender** option for a way of permitting non-trusted users to set envelope senders.

For a trusted user, there is never any check on the contents of the *From:* header line, and a *Sender:* line is never added. Furthermore, any existing *Sender:* line in incoming local (non-TCP/IP) messages is not removed.

Trusted users may also specify a host name, host address, interface address, protocol name, ident value, and authentication data when submitting a message locally. Thus, they are able to insert messages into Exim's queue locally that have the characteristics of messages received from a remote host. Untrusted users may in some circumstances use **-f**, but can never set the other values that are available to trusted users.

- The admin users are root, the Exim user, and any user that is a member of the Exim group or of any group listed in the **admin\_groups** configuration option. The current group does not have to be one of these groups.

Admin users are permitted to list the queue, and to carry out certain operations on messages, for example, to force delivery failures. It is also necessary to be an admin user in order to see the full information provided by the Exim monitor, and full debugging output.

By default, the use of the **-M**, **-q**, **-R**, and **-S** options to cause Exim to attempt delivery of messages on its queue is restricted to admin users. However, this restriction can be relaxed by setting the **prod\_requires\_admin** option false (that is, specifying **no\_prod\_requires\_admin**).

Similarly, the use of the **-bp** option to list all the messages in the queue is restricted to admin users unless **queue\_list\_requires\_admin** is set false.

**Warning:** If you configure your system so that admin users are able to edit Exim's configuration file, you are giving those users an easy way of getting root. There is further discussion of this issue at the start of chapter 6.

## 5.3 Command line options

Exim's command line options are described in alphabetical order below. If none of the options that specifies a specific action (such as starting the daemon or a queue runner, or testing an address, or receiving a message in a specific format, or listing the queue) are present, and there is at least one argument on the command line, **-bm** (accept a local message on the standard input, with the arguments specifying the recipients) is assumed. Otherwise, Exim outputs a brief message about itself and exits.

--

This is a pseudo-option whose only purpose is to terminate the options and therefore to cause subsequent command line items to be treated as arguments rather than options, even if they begin with hyphens.

### --help

This option causes Exim to output a few sentences stating what it is. The same output is generated if the Exim binary is called with no options and no arguments.

### --version

This option is an alias for **-bV** and causes version information to be displayed.

### -Ac

### -Am

These options are used by Sendmail for selecting configuration files and are ignored by Exim.

### -B<type>

This is a Sendmail option for selecting 7 or 8 bit processing. Exim is 8-bit clean; it ignores this option.

### -bd

This option runs Exim as a daemon, awaiting incoming SMTP connections. Usually the **-bd** option is combined with the **-q<time>** option, to specify that the daemon should also initiate periodic queue runs.

The **-bd** option can be used only by an admin user. If either of the **-d** (debugging) or **-v** (verifying) options are set, the daemon does not disconnect from the controlling terminal. When running this way, it can be stopped by pressing ctrl-C.

By default, Exim listens for incoming connections to the standard SMTP port on all the host's running interfaces. However, it is possible to listen on other ports, on multiple ports, and only on specific interfaces. Chapter 13 contains a description of the options that control this.

When a listening daemon is started without the use of **-oX** (that is, without overriding the normal configuration), it writes its process id to a file called *exim-daemon.pid* in Exim's spool directory. This location can be overridden by setting **PID\_FILE\_PATH** in *Local/Makefile*. The file is written while Exim is still running as root.

When **-oX** is used on the command line to start a listening daemon, the process id is not written to the normal pid file path. However, **-oP** can be used to specify a path on the command line if a pid file is required.

The SIGHUP signal can be used to cause the daemon to re-execute itself. This should be done whenever Exim's configuration file, or any file that is incorporated into it by means of the **.include** facility, is changed, and also whenever a new version of Exim is installed. It is not necessary to do this when other files that are referenced from the configuration (for example, alias files) are changed, because these are reread each time they are used.

#### **-bdf**

This option has the same effect as **-bd** except that it never disconnects from the controlling terminal, even when no debugging is specified.

#### **-be**

Run Exim in expansion testing mode. Exim discards its root privilege, to prevent ordinary users from using this mode to read otherwise inaccessible files. If no arguments are given, Exim runs interactively, prompting for lines of data. Otherwise, it processes each argument in turn.

If Exim was built with USE\_READLINE=yes in *Local/Makefile*, it tries to load the **libreadline** library dynamically whenever the **-be** option is used without command line arguments. If successful, it uses the *readline()* function, which provides extensive line-editing facilities, for reading the test data. A line history is supported.

Long expansion expressions can be split over several lines by using backslash continuations. As in Exim's run time configuration, white space at the start of continuation lines is ignored. Each argument or data line is passed through the string expansion mechanism, and the result is output. Variable values from the configuration file (for example, *\$qualify\_domain*) are available, but no message-specific values (such as *\$message\_exim\_id*) are set, because no message is being processed (but see **-bem** and **-Mset**).

**Note:** If you use this mechanism to test lookups, and you change the data files or databases you are using, you must exit and restart Exim before trying the same lookup again. Otherwise, because each Exim process caches the results of lookups, you will just get the same result as before.

Macro processing is done on lines before string-expansion: new macros can be defined and macros will be expanded. Because macros in the config file are often used for secrets, those are only available to admin users.

#### **-bem <filename>**

This option operates like **-be** except that it must be followed by the name of a file. For example:

```
exim -bem /tmp/testmessage
```

The file is read as a message (as if receiving a locally-submitted non-SMTP message) before any of the test expansions are done. Thus, message-specific variables such as *\$message\_size* and *\$header\_from:* are available. However, no *Received:* header is added to the message. If the **-t** option is set, recipients are read from the headers in the normal way, and are shown in the *\$recipients* variable. Note that recipients cannot be given on the command line, because further arguments are taken as strings to expand (just like **-be**).

#### **-bF <filename>**

This option is the same as **-bf** except that it assumes that the filter being tested is a system filter. The additional commands that are available only in system filters are recognized.

#### **-bf <filename>**

This option runs Exim in user filter testing mode; the file is the filter file to be tested, and a test message must be supplied on the standard input. If there are no message-dependent tests in the filter, an empty file can be supplied.

If you want to test a system filter file, use **-bF** instead of **-bf**. You can use both **-bF** and **-bf** on the same command, in order to test a system filter and a user filter in the same run. For example:

```
exim -bF /system/filter -bf /user/filter </test/message
```

This is helpful when the system filter adds header lines or sets filter variables that are used by the user filter.

If the test filter file does not begin with one of the special lines

```
# Exim filter
# Sieve filter
```

it is taken to be a normal *.forward* file, and is tested for validity under that interpretation. See sections 22.4 to 22.6 for a description of the possible contents of non-filter redirection lists.

The result of an Exim command that uses **-bf**, provided no errors are detected, is a list of the actions that Exim would try to take if presented with the message for real. More details of filter testing are given in the separate document entitled *Exim's interfaces to mail filtering*.

When testing a filter file, the envelope sender can be set by the **-f** option, or by a "From " line at the start of the test message. Various parameters that would normally be taken from the envelope recipient address of the message can be set by means of additional command line options (see the next four options).

**-bfd** <domain>

This sets the domain of the recipient address when a filter file is being tested by means of the **-bf** option. The default is the value of *\$qualify\_domain*.

**-bfl** <local part>

This sets the local part of the recipient address when a filter file is being tested by means of the **-bf** option. The default is the username of the process that calls Exim. A local part should be specified with any prefix or suffix stripped, because that is how it appears to the filter when a message is actually being delivered.

**-bfp** <prefix>

This sets the prefix of the local part of the recipient address when a filter file is being tested by means of the **-bf** option. The default is an empty prefix.

**-bfs** <suffix>

This sets the suffix of the local part of the recipient address when a filter file is being tested by means of the **-bf** option. The default is an empty suffix.

**-bh** <IP address>

This option runs a fake SMTP session as if from the given IP address, using the standard input and output. The IP address may include a port number at the end, after a full stop. For example:

```
exim -bh 10.9.8.7.1234
exim -bh fe80::a00:20ff:fe86:a061.5678
```

When an IPv6 address is given, it is converted into canonical form. In the case of the second example above, the value of *\$sender\_host\_address* after conversion to the canonical form is *fe80:0000:0000:0a00:20ff:fe86:a061.5678*.

Comments as to what is going on are written to the standard error file. These include lines beginning with "LOG" for anything that would have been logged. This facility is provided for testing configuration options for incoming messages, to make sure they implement the required policy. For example, you can test your relay controls using **-bh**.

**Warning 1:** You can test features of the configuration that rely on ident (RFC 1413) information by using the **-oMt** option. However, Exim cannot actually perform an ident callout when testing using **-bh** because there is no incoming SMTP connection.

**Warning 2:** Address verification callouts (see section 43.45) are also skipped when testing using **-bh**. If you want these callouts to occur, use **-bhc** instead.

Messages supplied during the testing session are discarded, and nothing is written to any of the real log files. There may be pauses when DNS (and other) lookups are taking place, and of course these may time out. The **-oMi** option can be used to specify a specific IP interface and port if this is important, and **-oMaa** and **-oMai** can be used to set parameters as if the SMTP session were authenticated.

The *exim\_checkaccess* utility is a “packaged” version of **-bh** whose output just states whether a given recipient address from a given host is acceptable or not. See section 53.8.

Features such as authentication and encryption, where the client input is not plain text, cannot easily be tested with **-bh**. Instead, you should use a specialized SMTP test program such as **swaks** (<http://jetmore.org/john/code/#swaks>).

**-bhc** <IP address>

This option operates in the same way as **-bh**, except that address verification callouts are performed if required. This includes consulting and updating the callout cache database.

**-bi**

Sendmail interprets the **-bi** option as a request to rebuild its alias file. Exim does not have the concept of a single alias file, and so it cannot mimic this behaviour. However, calls to */usr/lib/sendmail* with the **-bi** option tend to appear in various scripts such as NIS make files, so the option must be recognized.

If **-bi** is encountered, the command specified by the **bi\_command** configuration option is run, under the uid and gid of the caller of Exim. If the **-oA** option is used, its value is passed to the command as an argument. The command set by **bi\_command** may not contain arguments. The command can use the *exim\_dbmbuild* utility, or some other means, to rebuild alias files if this is required. If the **bi\_command** option is not set, calling Exim with **-bi** is a no-op.

**-bI:help**

We shall provide various options starting **-bI:** for querying Exim for information. The output of many of these will be intended for machine consumption. This one is not. The **-bI:help** option asks Exim for a synopsis of supported options beginning **-bI:**. Use of any of these options shall cause Exim to exit after producing the requested output.

**-bI:dscp**

This option causes Exim to emit an alphabetically sorted list of all recognised DSCP names.

**-bI:sieve**

This option causes Exim to emit an alphabetically sorted list of all supported Sieve protocol extensions on stdout, one per line. This is anticipated to be useful for ManageSieve (RFC 5804) implementations, in providing that protocol’s SIEVE capability response line. As the precise list may depend upon compile-time build options, which this option will adapt to, this is the only way to guarantee a correct response.

**-bm**

This option runs an Exim receiving process that accepts an incoming, locally-generated message on the standard input. The recipients are given as the command arguments (except when **-t** is also present – see below). Each argument can be a comma-separated list of RFC 2822 addresses. This is the default option for selecting the overall action of an Exim call; it is assumed if no other conflicting option is present.

If any addresses in the message are unqualified (have no domain), they are qualified by the values of the **qualify\_domain** or **qualify\_recipient** options, as appropriate. The **-bnq** option (see below) provides a way of suppressing this for special cases.

Policy checks on the contents of local messages can be enforced by means of the non-SMTP ACL. See chapter 43 for details.

The return code is zero if the message is successfully accepted. Otherwise, the action is controlled by the **-oex** option setting – see below.

The format of the message must be as defined in RFC 2822, except that, for compatibility with Sendmail and Smail, a line in one of the forms

```
From sender Fri Jan 5 12:55 GMT 1997
From sender Fri, 5 Jan 97 12:55:01
```

(with the weekday optional, and possibly with additional text after the date) is permitted to appear at the start of the message. There appears to be no authoritative specification of the format of this

line. Exim recognizes it by matching against the regular expression defined by the **uucp\_from\_pattern** option, which can be changed if necessary.

The specified sender is treated as if it were given as the argument to the **-f** option, but if a **-f** option is also present, its argument is used in preference to the address taken from the message. The caller of Exim must be a trusted user for the sender of a message to be set in this way.

### **-bmalware** <filename>

This debugging option causes Exim to scan the given file or directory (depending on the used scanner interface), using the malware scanning framework. The option of **av\_scanner** influences this option, so if **av\_scanner**'s value is dependent upon an expansion then the expansion should have defaults which apply to this invocation. ACLs are not invoked, so if **av\_scanner** references an ACL variable then that variable will never be populated and **-bmalware** will fail.

Exim will have changed working directory before resolving the filename, so using fully qualified pathnames is advisable. Exim will be running as the Exim user when it tries to open the file, rather than as the invoking user. This option requires admin privileges.

The **-bmalware** option will not be extended to be more generally useful, there are better tools for file-scanning. This option exists to help administrators verify their Exim and AV scanner configuration.

### **-bnq**

By default, Exim automatically qualifies unqualified addresses (those without domains) that appear in messages that are submitted locally (that is, not over TCP/IP). This qualification applies both to addresses in envelopes, and addresses in header lines. Sender addresses are qualified using **qualify\_domain**, and recipient addresses using **qualify\_recipient** (which defaults to the value of **qualify\_domain**).

Sometimes, qualification is not wanted. For example, if **-bS** (batch SMTP) is being used to re-submit messages that originally came from remote hosts after content scanning, you probably do not want to qualify unqualified addresses in header lines. (Such lines will be present only if you have not enabled a header syntax check in the appropriate ACL.)

The **-bnq** option suppresses all qualification of unqualified addresses in messages that originate on the local host. When this is used, unqualified addresses in the envelope provoke errors (causing message rejection) and unqualified addresses in header lines are left alone.

### **-bP**

If this option is given with no arguments, it causes the values of all Exim's main configuration options to be written to the standard output. The values of one or more specific options can be requested by giving their names as arguments, for example:

```
exim -bP qualify_domain hold_domains
```

However, any option setting that is preceded by the word "hide" in the configuration file is not shown in full, except to an admin user. For other users, the output is as in this example:

```
mysql_servers = <value not displayable>
```

If **config** is given as an argument, the config is output, as it was parsed, any include file resolved, any comment removed.

If **config\_file** is given as an argument, the name of the run time configuration file is output. (**configure\_file** works too, for backward compatibility.) If a list of configuration files was supplied, the value that is output here is the name of the file that was actually used.

If the **-n** flag is given, then for most modes of **-bP** operation the name will not be output.

If **log\_file\_path** or **pid\_file\_path** are given, the names of the directories where log files and daemon pid files are written are output, respectively. If these values are unset, log files are written in a sub-directory of the spool directory called **log**, and the pid file is written directly into the spool directory.

If **-bP** is followed by a name preceded by +, for example,

```
exim -bP +local_domains
```

it searches for a matching named list of any type (domain, host, address, or local part) and outputs what it finds.

If one of the words **router**, **transport**, or **authenticator** is given, followed by the name of an appropriate driver instance, the option settings for that driver are output. For example:

```
exim -bP transport local_delivery
```

The generic driver options are output first, followed by the driver's private options. A list of the names of drivers of a particular type can be obtained by using one of the words **router\_list**, **transport\_list**, or **authenticator\_list**, and a complete list of all drivers with their option settings can be obtained by using **routers**, **transports**, or **authenticators**.

If **environment** is given as an argument, the set of environment variables is output, line by line. Using the **-n** flag suppresses the value of the variables.

If invoked by an admin user, then **macro**, **macro\_list** and **macros** are available, similarly to the drivers. Because macros are sometimes used for storing passwords, this option is restricted. The output format is one item per line.

### **-bp**

This option requests a listing of the contents of the mail queue on the standard output. If the **-bp** option is followed by a list of message ids, just those messages are listed. By default, this option can be used only by an admin user. However, the **queue\_list\_requires\_admin** option can be set false to allow any user to see the queue.

Each message on the queue is displayed as in the following example:

```
25m 2.9K 0t5C6f-0000c8-00 <alice@wonderland.fict.example>
red.king@looking-glass.fict.example
<other addresses>
```

The first line contains the length of time the message has been on the queue (in this case 25 minutes), the size of the message (2.9K), the unique local identifier for the message, and the message sender, as contained in the envelope. For bounce messages, the sender address is empty, and appears as "<>". If the message was submitted locally by an untrusted user who overrode the default sender address, the user's login name is shown in parentheses before the sender address.

If the message is frozen (attempts to deliver it are suspended) then the text "\*\*\*\* frozen \*\*\*\*" is displayed at the end of this line.

The recipients of the message (taken from the envelope, not the headers) are displayed on subsequent lines. Those addresses to which the message has already been delivered are marked with the letter D. If an original address gets expanded into several addresses via an alias or forward file, the original is displayed with a D only when deliveries for all of its child addresses are complete.

### **-bpa**

This option operates like **-bp**, but in addition it shows delivered addresses that were generated from the original top level address(es) in each message by alias or forwarding operations. These addresses are flagged with "+D" instead of just "D".

### **-bpc**

This option counts the number of messages on the queue, and writes the total to the standard output. It is restricted to admin users, unless **queue\_list\_requires\_admin** is set false.

### **-bpr**

This option operates like **-bp**, but the output is not sorted into chronological order of message arrival. This can speed it up when there are lots of messages on the queue, and is particularly useful if the output is going to be post-processed in a way that doesn't need the sorting.

### **-bptra**

This option is a combination of **-bpr** and **-bpa**.



### **-bpru**

This option is a combination of **-bpr** and **-bpu**.

### **-bpu**

This option operates like **-bp** but shows only undelivered top-level addresses for each message displayed. Addresses generated by aliasing or forwarding are not shown, unless the message was deferred after processing by a router with the **one\_time** option set.

### **-brt**

This option is for testing retry rules, and it must be followed by up to three arguments. It causes Exim to look for a retry rule that matches the values and to write it to the standard output. For example:

```
exim -brt bach.comp.mus.example
Retry rule: *.comp.mus.example F,2h,15m; F,4d,30m;
```

See chapter 32 for a description of Exim's retry rules. The first argument, which is required, can be a complete address in the form *local\_part@domain*, or it can be just a domain name. If the second argument contains a dot, it is interpreted as an optional second domain name; if no retry rule is found for the first argument, the second is tried. This ties in with Exim's behaviour when looking for retry rules for remote hosts – if no rule is found that matches the host, one that matches the mail domain is sought. Finally, an argument that is the name of a specific delivery error, as used in setting up retry rules, can be given. For example:

```
exim -brt haydn.comp.mus.example quota_3d
Retry rule: *@haydn.comp.mus.example quota_3d F,1h,15m
```

### **-brw**

This option is for testing address rewriting rules, and it must be followed by a single argument, consisting of either a local part without a domain, or a complete address with a fully qualified domain. Exim outputs how this address would be rewritten for each possible place it might appear. See chapter 31 for further details.

### **-bS**

This option is used for batched SMTP input, which is an alternative interface for non-interactive local message submission. A number of messages can be submitted in a single run. However, despite its name, this is not really SMTP input. Exim reads each message's envelope from SMTP commands on the standard input, but generates no responses. If the caller is trusted, or **untrusted\_set\_sender** is set, the senders in the SMTP MAIL commands are believed; otherwise the sender is always the caller of Exim.

The message itself is read from the standard input, in SMTP format (leading dots doubled), terminated by a line containing just a single dot. An error is provoked if the terminating dot is missing. A further message may then follow.

As for other local message submissions, the contents of incoming batch SMTP messages can be checked using the non-SMTP ACL (see chapter 43). Unqualified addresses are automatically qualified using **qualify\_domain** and **qualify\_recipient**, as appropriate, unless the **-bnq** option is used.

Some other SMTP commands are recognized in the input. HELO and EHLO act as RSET; VRFY, EXPN, ETRN, and HELP act as NOOP; QUIT quits, ignoring the rest of the standard input.

If any error is encountered, reports are written to the standard output and error streams, and Exim gives up immediately. The return code is 0 if no error was detected; it is 1 if one or more messages were accepted before the error was detected; otherwise it is 2.

More details of input using batched SMTP are given in section 48.11.

### **-bs**

This option causes Exim to accept one or more messages by reading SMTP commands on the standard input, and producing SMTP replies on the standard output. SMTP policy controls, as defined in ACLs (see chapter 43) are applied. Some user agents use this interface as a way of passing locally-generated messages to the MTA.

In this usage, if the caller of Exim is trusted, or **untrusted\_set\_sender** is set, the senders of messages are taken from the SMTP MAIL commands. Otherwise the content of these commands is ignored and the sender is set up as the calling user. Unqualified addresses are automatically qualified using **qualify\_domain** and **qualify\_recipient**, as appropriate, unless the **-bnq** option is used.

The **-bs** option is also used to run Exim from *inetd*, as an alternative to using a listening daemon. Exim can distinguish the two cases by checking whether the standard input is a TCP/IP socket. When Exim is called from *inetd*, the source of the mail is assumed to be remote, and the comments above concerning senders and qualification do not apply. In this situation, Exim behaves in exactly the same way as it does when receiving a message via the listening daemon.

#### **-bt**

This option runs Exim in address testing mode, in which each argument is taken as a recipient address to be tested for deliverability. The results are written to the standard output. If a test fails, and the caller is not an admin user, no details of the failure are output, because these might contain sensitive information such as usernames and passwords for database lookups.

If no arguments are given, Exim runs in an interactive manner, prompting with a right angle bracket for addresses to be tested.

Unlike the **-be** test option, you cannot arrange for Exim to use the *readline()* function, because it is running as *root* and there are security issues.

Each address is handled as if it were the recipient address of a message (compare the **-bv** option). It is passed to the routers and the result is written to the standard output. However, any router that has **no\_address\_test** set is bypassed. This can make **-bt** easier to use for genuine routing tests if your first router passes everything to a scanner program.

The return code is 2 if any address failed outright; it is 1 if no address failed outright but at least one could not be resolved for some reason. Return code 0 is given only when all addresses succeed.

**Note:** When actually delivering a message, Exim removes duplicate recipient addresses after routing is complete, so that only one delivery takes place. This does not happen when testing with **-bt**; the full results of routing are always shown.

**Warning:** **-bt** can only do relatively simple testing. If any of the routers in the configuration makes any tests on the sender address of a message, you can use the **-f** option to set an appropriate sender when running **-bt** tests. Without it, the sender is assumed to be the calling user at the default qualifying domain. However, if you have set up (for example) routers whose behaviour depends on the contents of an incoming message, you cannot test those conditions using **-bt**. The **-N** option provides a possible way of doing such tests.

#### **-bV**

This option causes Exim to write the current version number, compilation number, and compilation date of the *exim* binary to the standard output. It also lists the DBM library that is being used, the optional modules (such as specific lookup types), the drivers that are included in the binary, and the name of the run time configuration file that is in use.

As part of its operation, **-bV** causes Exim to read and syntax check its configuration file. However, this is a static check only. It cannot check values that are to be expanded. For example, although a misspelt ACL verb is detected, an error in the verb's arguments is not. You cannot rely on **-bV** alone to discover (for example) all the typos in the configuration; some realistic testing is needed. The **-bh** and **-N** options provide more dynamic testing facilities.

#### **-bv**

This option runs Exim in address verification mode, in which each argument is taken as a recipient address to be verified by the routers. (This does not involve any verification callouts). During normal operation, verification happens mostly as a consequence processing a **verify** condition in an ACL (see chapter 43). If you want to test an entire ACL, possibly including callouts, see the **-bh** and **-bhc** options.

If verification fails, and the caller is not an admin user, no details of the failure are output, because these might contain sensitive information such as usernames and passwords for database lookups.

If no arguments are given, Exim runs in an interactive manner, prompting with a right angle bracket for addresses to be verified.

Unlike the **-be** test option, you cannot arrange for Exim to use the *readline()* function, because it is running as *exim* and there are security issues.

Verification differs from address testing (the **-bt** option) in that routers that have **no\_verify** set are skipped, and if the address is accepted by a router that has **fail\_verify** set, verification fails. The address is verified as a recipient if **-bv** is used; to test verification for a sender address, **-bvs** should be used.

If the **-v** option is not set, the output consists of a single line for each address, stating whether it was verified or not, and giving a reason in the latter case. Without **-v**, generating more than one address by redirection causes verification to end successfully, without considering the generated addresses. However, if just one address is generated, processing continues, and the generated address must verify successfully for the overall verification to succeed.

When **-v** is set, more details are given of how the address has been handled, and in the case of address redirection, all the generated addresses are also considered. Verification may succeed for some and fail for others.

The return code is 2 if any address failed outright; it is 1 if no address failed outright but at least one could not be resolved for some reason. Return code 0 is given only when all addresses succeed.

If any of the routers in the configuration makes any tests on the sender address of a message, you should use the **-f** option to set an appropriate sender when running **-bv** tests. Without it, the sender is assumed to be the calling user at the default qualifying domain.

#### **-bvs**

This option acts like **-bv**, but verifies the address as a sender rather than a recipient address. This affects any rewriting and qualification that might happen.

#### **-bw**

This option runs Exim as a daemon, awaiting incoming SMTP connections, similarly to the **-bd** option. All port specifications on the command-line and in the configuration file are ignored. Queue-running may not be specified.

In this mode, Exim expects to be passed a socket as fd 0 (stdin) which is listening for connections. This permits the system to start up and have inetd (or equivalent) listen on the SMTP ports, starting an Exim daemon for each port only when the first connection is received.

If the option is given as **-bw<time>** then the time is a timeout, after which the daemon will exit, which should cause inetd to listen once more.

#### **-C <filelist>**

This option causes Exim to find the run time configuration file from the given list instead of from the list specified by the `CONFIGURE_FILE` compile-time setting. Usually, the list will consist of just a single file name, but it can be a colon-separated list of names. In this case, the first file that exists is used. Failure to open an existing file stops Exim from proceeding any further along the list, and an error is generated.

When this option is used by a caller other than root, and the list is different from the compiled-in list, Exim gives up its root privilege immediately, and runs with the real and effective uid and gid set to those of the caller. However, if a `TRUSTED_CONFIG_LIST` file is defined in *Local/Makefile*, that file contains a list of full pathnames, one per line, for configuration files which are trusted. Root privilege is retained for any configuration file so listed, as long as the caller is the Exim user (or the user specified in the `CONFIGURE_OWNER` option, if any), and as long as the configuration file is not writeable by inappropriate users or groups.

Leaving `TRUSTED_CONFIG_LIST` unset precludes the possibility of testing a configuration using **-C** right through message reception and delivery, even if the caller is root. The reception

works, but by that time, Exim is running as the Exim user, so when it re-executes to regain privilege for the delivery, the use of **-C** causes privilege to be lost. However, root can test reception and delivery using two separate commands (one to put a message on the queue, using **-odq**, and another to do the delivery, using **-M**).

If `ALT_CONFIG_PREFIX` is defined in *Local/Makefile*, it specifies a prefix string with which any file named in a **-C** command line option must start. In addition, the file name must not contain the sequence `/ . . /`. However, if the value of the **-C** option is identical to the value of `CONFIGURE_FILE` in *Local/Makefile*, Exim ignores **-C** and proceeds as usual. There is no default setting for `ALT_CONFIG_PREFIX`; when it is unset, any file name can be used with **-C**.

`ALT_CONFIG_PREFIX` can be used to confine alternative configuration files to a directory to which only root has access. This prevents someone who has broken into the Exim account from running a privileged Exim with an arbitrary configuration file.

The **-C** facility is useful for ensuring that configuration files are syntactically correct, but cannot be used for test deliveries, unless the caller is privileged, or unless it is an exotic configuration that does not require privilege. No check is made on the owner or group of the files specified by this option.

#### **-D**<macro>=<value>

This option can be used to override macro definitions in the configuration file (see section 6.4). However, like **-C**, if it is used by an unprivileged caller, it causes Exim to give up its root privilege. If `DISABLE_D_OPTION` is defined in *Local/Makefile*, the use of **-D** is completely disabled, and its use causes an immediate error exit.

If `WHITELIST_D_MACROS` is defined in *Local/Makefile* then it should be a colon-separated list of macros which are considered safe and, if **-D** only supplies macros from this list, and the values are acceptable, then Exim will not give up root privilege if the caller is root, the Exim run-time user, or the `CONFIGURE_OWNER`, if set. This is a transition mechanism and is expected to be removed in the future. Acceptable values for the macros satisfy the regexp: `^[A-Za-z0-9_/.-]*$`

The entire option (including equals sign if present) must all be within one command line item. **-D** can be used to set the value of a macro to the empty string, in which case the equals sign is optional. These two commands are synonymous:

```
exim -DABC ...
exim -DABC= ...
```

To include spaces in a macro definition item, quotes must be used. If you use quotes, spaces are permitted around the macro name and the equals sign. For example:

```
exim '-D ABC = something' ...
```

**-D** may be repeated up to 10 times on a command line. Only macro names up to 22 letters long can be set.

#### **-d**<debug options>

This option causes debugging information to be written to the standard error stream. It is restricted to admin users because debugging output may show database queries that contain password information. Also, the details of users' filter files should be protected. If a non-admin user uses **-d**, Exim writes an error message to the standard error stream and exits with a non-zero return code.

When **-d** is used, **-v** is assumed. If **-d** is given on its own, a lot of standard debugging data is output. This can be reduced, or increased to include some more rarely needed information, by directly following **-d** with a string made up of names preceded by plus or minus characters. These add or remove sets of debugging data, respectively. For example, **-d+filter** adds filter debugging, whereas **-d-all+filter** selects only filter debugging. Note that no spaces are allowed in the debug setting. The available debugging categories are:

<code>acl</code>	ACL interpretation
<code>auth</code>	authenticators
<code>deliver</code>	general delivery logic

<code>dns</code>	DNS lookups (see also <code>resolver</code> )
<code>dnsbl</code>	DNS black list (aka RBL) code
<code>exec</code>	arguments for <code>execv()</code> calls
<code>expand</code>	detailed debugging for string expansions
<code>filter</code>	filter handling
<code>hints_lookup</code>	hints data lookups
<code>host_lookup</code>	all types of name-to-IP address handling
<code>ident</code>	ident lookup
<code>interface</code>	lists of local interfaces
<code>lists</code>	matching things in lists
<code>load</code>	system load checks
<code>local_scan</code>	can be used by <code>local_scan()</code> (see chapter 45)
<code>lookup</code>	general lookup code and all lookups
<code>memory</code>	memory handling
<code>pid</code>	add <code>pid</code> to debug output lines
<code>process_info</code>	setting info for the process log
<code>queue_run</code>	queue runs
<code>receive</code>	general message reception logic
<code>resolver</code>	turn on the DNS resolver's debugging output
<code>retry</code>	retry handling
<code>rewrite</code>	address rewriting
<code>route</code>	address routing
<code>timestamp</code>	add timestamp to debug output lines
<code>tls</code>	TLS logic
<code>transport</code>	transports
<code>uid</code>	changes of <code>uid/gid</code> and looking up <code>uid/gid</code>
<code>verify</code>	address verification logic
<code>all</code>	almost all of the above (see below), and also <code>-v</code>

The `all` option excludes `memory` when used as `+all`, but includes it for `-all`. The reason for this is that `+all` is something that people tend to use when generating debug output for Exim maintainers. If `+memory` is included, an awful lot of output that is very rarely of interest is generated, so it now has to be explicitly requested. However, `-all` does turn everything off.

The `resolver` option produces output only if the DNS resolver was compiled with `DEBUG` enabled. This is not the case in some operating systems. Also, unfortunately, debugging output from the DNS resolver is written to `stdout` rather than `stderr`.

The default (`-d` with no argument) omits `expand`, `filter`, `interface`, `load`, `memory`, `pid`, `resolver`, and `timestamp`. However, the `pid` selector is forced when debugging is turned on for a daemon, which then passes it on to any re-executed Exims. Exim also automatically adds the `pid` to debug lines when several remote deliveries are run in parallel.

The `timestamp` selector causes the current time to be inserted at the start of all debug output lines. This can be useful when trying to track down delays in processing.

If the `debug_print` option is set in any driver, it produces output whenever any debugging is selected, or if `-v` is used.

#### **-dd**<debug options>

This option behaves exactly like `-d` except when used on a command that starts a daemon process. In that case, debugging is turned off for the subprocesses that the daemon creates. Thus, it is useful for monitoring the behaviour of the daemon without creating as much output as full debugging does.

#### **-droper**

This is an obsolete option that is now a no-op. It used to affect the way Exim handled CR and LF characters in incoming messages. What happens now is described in section 47.2.

#### **-E**

This option specifies that an incoming message is a locally-generated delivery failure report. It is used internally by Exim when handling delivery failures and is not intended for external use. Its

only effect is to stop Exim generating certain messages to the postmaster, as otherwise message cascades could occur in some situations. As part of the same option, a message id may follow the characters **-E**. If it does, the log entry for the receipt of the new message contains the id, following “R=”, as a cross-reference.

#### **-ex**

There are a number of Sendmail options starting with **-oe** which seem to be called by various programs without the leading **o** in the option. For example, the **vacation** program uses **-eq**. Exim treats all options of the form **-ex** as synonymous with the corresponding **-oex** options.

#### **-F** <string>

This option sets the sender’s full name for use when a locally-generated message is being accepted. In the absence of this option, the user’s *gecos* entry from the password data is used. As users are generally permitted to alter their *gecos* entries, no security considerations are involved. White space between **-F** and the <string> is optional.

#### **-f** <address>

This option sets the address of the envelope sender of a locally-generated message (also known as the return path). The option can normally be used only by a trusted user, but **untrusted\_set\_sender** can be set to allow untrusted users to use it.

Processes running as root or the Exim user are always trusted. Other trusted users are defined by the **trusted\_users** or **trusted\_groups** options. In the absence of **-f**, or if the caller is not trusted, the sender of a local message is set to the caller’s login name at the default qualify domain.

There is one exception to the restriction on the use of **-f**: an empty sender can be specified by any user, trusted or not, to create a message that can never provoke a bounce. An empty sender can be specified either as an empty string, or as a pair of angle brackets with nothing between them, as in these examples of shell commands:

```
exim -f '<>' user@domain
exim -f "" user@domain
```

In addition, the use of **-f** is not restricted when testing a filter file with **-bf** or when testing or verifying addresses using the **-bt** or **-bv** options.

Allowing untrusted users to change the sender address does not of itself make it possible to send anonymous mail. Exim still checks that the *From:* header refers to the local user, and if it does not, it adds a *Sender:* header, though this can be overridden by setting **no\_local\_from\_check**.

White space between **-f** and the <address> is optional (that is, they can be given as two arguments or one combined argument). The sender of a locally-generated message can also be set (when permitted) by an initial “From ” line in the message – see the description of **-bm** above – but if **-f** is also present, it overrides “From ”.

#### **-G**

This option is equivalent to an ACL applying:

```
control = suppress_local_fixups
```

for every message received. Note that Sendmail will complain about such bad formatting, where Exim silently just does not fix it up. This may change in future.

As this affects audit information, the caller must be a trusted user to use this option.

#### **-h** <number>

This option is accepted for compatibility with Sendmail, but has no effect. (In Sendmail it overrides the “hop count” obtained by counting *Received:* headers.)

#### **-i**

This option, which has the same effect as **-oi**, specifies that a dot on a line by itself should not terminate an incoming, non-SMTP message. I can find no documentation for this option in Solaris 2.4 Sendmail, but the *mailx* command in Solaris 2.4 uses it. See also **-ti**.

**-L** *<tag>*

This option is equivalent to setting **syslog\_processname** in the config file and setting **log\_file\_path** to **syslog**. Its use is restricted to administrators. The configuration file has to be read and parsed, to determine access rights, before this is set and takes effect, so early configuration file errors will not honour this flag.

The tag should not be longer than 32 characters.

**-M** *<message id> <message id> ...*

This option requests Exim to run a delivery attempt on each message in turn. If any of the messages are frozen, they are automatically thawed before the delivery attempt. The settings of **queue\_domains**, **queue\_smtp\_domains**, and **hold\_domains** are ignored.

Retry hints for any of the addresses are overridden – Exim tries to deliver even if the normal retry time has not yet been reached. This option requires the caller to be an admin user. However, there is an option called **prod\_requires\_admin** which can be set false to relax this restriction (and also the same requirement for the **-q**, **-R**, and **-S** options).

The deliveries happen synchronously, that is, the original Exim process does not terminate until all the delivery attempts have finished. No output is produced unless there is a serious error. If you want to see what is happening, use the **-v** option as well, or inspect Exim's main log.

**-Mar** *<message id> <address> <address> ...*

This option requests Exim to add the addresses to the list of recipients of the message (“ar” for “add recipients”). The first argument must be a message id, and the remaining ones must be email addresses. However, if the message is active (in the middle of a delivery attempt), it is not altered. This option can be used only by an admin user.

**-MC** *<transport> <hostname> <sequence number> <message id>*

This option is not intended for use by external callers. It is used internally by Exim to invoke another instance of itself to deliver a waiting message using an existing SMTP connection, which is passed as the standard input. Details are given in chapter 48. This must be the final option, and the caller must be root or the Exim user in order to use it.

**-MCA**

This option is not intended for use by external callers. It is used internally by Exim in conjunction with the **-MC** option. It signifies that the connection to the remote host has been authenticated.

**-MCD**

This option is not intended for use by external callers. It is used internally by Exim in conjunction with the **-MC** option. It signifies that the remote host supports the ESMTP *DSN* extension.

**-MCG** *<queue name>*

This option is not intended for use by external callers. It is used internally by Exim in conjunction with the **-MC** option. It signifies that an alternate queue is used, named by the following argument.

**-MCK**

This option is not intended for use by external callers. It is used internally by Exim in conjunction with the **-MC** option. It signifies that a remote host supports the ESMTP *CHUNKING* extension.

**-MCP**

This option is not intended for use by external callers. It is used internally by Exim in conjunction with the **-MC** option. It signifies that the server to which Exim is connected supports pipelining.

**-MCQ** *<process id> <pipe fd>*

This option is not intended for use by external callers. It is used internally by Exim in conjunction with the **-MC** option when the original delivery was started by a queue runner. It passes on the process id of the queue runner, together with the file descriptor number of an open pipe. Closure of the pipe signals the final completion of the sequence of processes that are passing messages through the same SMTP connection.

## **-MCS**

This option is not intended for use by external callers. It is used internally by Exim in conjunction with the **-MC** option, and passes on the fact that the SMTP SIZE option should be used on messages delivered down the existing connection.

## **-MCT**

This option is not intended for use by external callers. It is used internally by Exim in conjunction with the **-MC** option, and passes on the fact that the host to which Exim is connected supports TLS encryption.

## **-MCt** <IP address> <port> <cipher>

This option is not intended for use by external callers. It is used internally by Exim in conjunction with the **-MC** option, and passes on the fact that the connection is being proxied by a parent process for handling TLS encryption. The arguments give the local address and port being proxied, and the TLS cipher.

## **-Mc** <message id> <message id> ...

This option requests Exim to run a delivery attempt on each message in turn, but unlike the **-M** option, it does check for retry hints, and respects any that are found. This option is not very useful to external callers. It is provided mainly for internal use by Exim when it needs to re-invoke itself in order to regain root privilege for a delivery (see chapter 55). However, **-Mc** can be useful when testing, in order to run a delivery that respects retry times and other options such as **hold\_domains** that are overridden when **-M** is used. Such a delivery does not count as a queue run. If you want to run a specific delivery as if in a queue run, you should use **-q** with a message id argument. A distinction between queue run deliveries and other deliveries is made in one or two places.

## **-Mes** <message id> <address>

This option requests Exim to change the sender address in the message to the given address, which must be a fully qualified address or “<>” (“es” for “edit sender”). There must be exactly two arguments. The first argument must be a message id, and the second one an email address. However, if the message is active (in the middle of a delivery attempt), its status is not altered. This option can be used only by an admin user.

## **-Mf** <message id> <message id> ...

This option requests Exim to mark each listed message as “frozen”. This prevents any delivery attempts taking place until the message is “thawed”, either manually or as a result of the **auto\_thaw** configuration option. However, if any of the messages are active (in the middle of a delivery attempt), their status is not altered. This option can be used only by an admin user.

## **-Mg** <message id> <message id> ...

This option requests Exim to give up trying to deliver the listed messages, including any that are frozen. However, if any of the messages are active, their status is not altered. For non-bounce messages, a delivery error message is sent to the sender, containing the text “cancelled by administrator”. Bounce messages are just discarded. This option can be used only by an admin user.

## **-Mmad** <message id> <message id> ...

This option requests Exim to mark all the recipient addresses in the messages as already delivered (“mad” for “mark all delivered”). However, if any message is active (in the middle of a delivery attempt), its status is not altered. This option can be used only by an admin user.

## **-Mmd** <message id> <address> <address> ...

This option requests Exim to mark the given addresses as already delivered (“md” for “mark delivered”). The first argument must be a message id, and the remaining ones must be email addresses. These are matched to recipient addresses in the message in a case-sensitive manner. If the message is active (in the middle of a delivery attempt), its status is not altered. This option can be used only by an admin user.

## **-Mrm** <message id> <message id> ...

This option requests Exim to remove the given messages from the queue. No bounce messages are sent; each message is simply forgotten. However, if any of the messages are active, their status is not altered. This option can be used only by an admin user or by the user who originally caused the message to be placed on the queue.



**-Mset** <message id>

This option is useful only in conjunction with **-be** (that is, when testing string expansions). Exim loads the given message from its spool before doing the test expansions, thus setting message-specific variables such as *\$message\_size* and the header variables. The *\$recipients* variable is made available. This feature is provided to make it easier to test expansions that make use of these variables. However, this option can be used only by an admin user. See also **-bem**.

**-Mt** <message id> <message id> ...

This option requests Exim to “thaw” any of the listed messages that are “frozen”, so that delivery attempts can resume. However, if any of the messages are active, their status is not altered. This option can be used only by an admin user.

**-Mvb** <message id>

This option causes the contents of the message body (-D) spool file to be written to the standard output. This option can be used only by an admin user.

**-Mvc** <message id>

This option causes a copy of the complete message (header lines plus body) to be written to the standard output in RFC 2822 format. This option can be used only by an admin user.

**-Mvh** <message id>

This option causes the contents of the message headers (-H) spool file to be written to the standard output. This option can be used only by an admin user.

**-Mvl** <message id>

This option causes the contents of the message log spool file to be written to the standard output. This option can be used only by an admin user.

**-m**

This is apparently a synonym for **-om** that is accepted by Sendmail, so Exim treats it that way too.

**-N**

This is a debugging option that inhibits delivery of a message at the transport level. It implies **-v**. Exim goes through many of the motions of delivery – it just doesn’t actually transport the message, but instead behaves as if it had successfully done so. However, it does not make any updates to the retry database, and the log entries for deliveries are flagged with “\*>” rather than “=>”.

Because **-N** discards any message to which it applies, only root or the Exim user are allowed to use it with **-bd**, **-q**, **-R** or **-M**. In other words, an ordinary user can use it only when supplying an incoming message to which it will apply. Although transportation never fails when **-N** is set, an address may be deferred because of a configuration problem on a transport, or a routing problem. Once **-N** has been used for a delivery attempt, it sticks to the message, and applies to any subsequent delivery attempts that may happen for that message.

**-n**

This option is interpreted by Sendmail to mean “no aliasing”. For normal modes of operation, it is ignored by Exim. When combined with **-bP** it makes the output more terse (suppresses option names, environment values and config pretty printing).

**-O** <data>

This option is interpreted by Sendmail to mean `set option`. It is ignored by Exim.

**-oA** <file name>

This option is used by Sendmail in conjunction with **-bi** to specify an alternative alias file name. Exim handles **-bi** differently; see the description above.

**-oB** <n>

This is a debugging option which limits the maximum number of messages that can be delivered down one SMTP connection, overriding the value set in any *smtp* transport. If <n> is omitted, the limit is set to 1.

**-odb**

This option applies to all modes in which Exim accepts incoming messages, including the listening daemon. It requests “background” delivery of such messages, which means that the accepting

process automatically starts a delivery process for each message received, but does not wait for the delivery processes to finish.

When all the messages have been received, the reception process exits, leaving the delivery processes to finish in their own time. The standard output and error streams are closed at the start of each delivery process. This is the default action if none of the **-od** options are present.

If one of the queueing options in the configuration file (**queue\_only** or **queue\_only\_file**, for example) is in effect, **-odb** overrides it if **queue\_only\_override** is set true, which is the default setting. If **queue\_only\_override** is set false, **-odb** has no effect.

#### **-odf**

This option requests “foreground” (synchronous) delivery when Exim has accepted a locally-generated message. (For the daemon it is exactly the same as **-odb**.) A delivery process is automatically started to deliver the message, and Exim waits for it to complete before proceeding.

The original Exim reception process does not finish until the delivery process for the final message has ended. The standard error stream is left open during deliveries.

However, like **-odb**, this option has no effect if **queue\_only\_override** is false and one of the queueing options in the configuration file is in effect.

If there is a temporary delivery error during foreground delivery, the message is left on the queue for later delivery, and the original reception process exits. See chapter 51 for a way of setting up a restricted configuration that never queues messages.

#### **-odi**

This option is synonymous with **-odf**. It is provided for compatibility with Sendmail.

#### **-odq**

This option applies to all modes in which Exim accepts incoming messages, including the listening daemon. It specifies that the accepting process should not automatically start a delivery process for each message received. Messages are placed on the queue, and remain there until a subsequent queue runner process encounters them. There are several configuration options (such as **queue\_only**) that can be used to queue incoming messages under certain conditions. This option overrides all of them and also **-odqs**. It always forces queueing.

#### **-odqs**

This option is a hybrid between **-odb/-odi** and **-odq**. However, like **-odb** and **-odi**, this option has no effect if **queue\_only\_override** is false and one of the queueing options in the configuration file is in effect.

When **-odqs** does operate, a delivery process is started for each incoming message, in the background by default, but in the foreground if **-odi** is also present. The recipient addresses are routed, and local deliveries are done in the normal way. However, if any SMTP deliveries are required, they are not done at this time, so the message remains on the queue until a subsequent queue runner process encounters it. Because routing was done, Exim knows which messages are waiting for which hosts, and so a number of messages for the same host can be sent in a single SMTP connection. The **queue\_smtp\_domains** configuration option has the same effect for specific domains. See also the **-qq** option.

#### **-oee**

If an error is detected while a non-SMTP message is being received (for example, a malformed address), the error is reported to the sender in a mail message.

Provided this error message is successfully sent, the Exim receiving process exits with a return code of zero. If not, the return code is 2 if the problem is that the original message has no recipients, or 1 for any other error. This is the default **-oex** option if Exim is called as *rmail*.

#### **-oem**

This is the same as **-oee**, except that Exim always exits with a non-zero return code, whether or not the error message was successfully sent. This is the default **-oex** option, unless Exim is called as *rmail*.

**-oep**

If an error is detected while a non-SMTP message is being received, the error is reported by writing a message to the standard error file (stderr). The return code is 1 for all errors.

**-oeq**

This option is supported for compatibility with Sendmail, but has the same effect as **-oep**.

**-oew**

This option is supported for compatibility with Sendmail, but has the same effect as **-oem**.

**-oi**

This option, which has the same effect as **-i**, specifies that a dot on a line by itself should not terminate an incoming, non-SMTP message. Otherwise, a single dot does terminate, though Exim does no special processing for other lines that start with a dot. This option is set by default if Exim is called as *rmail*. See also **-ti**.

**-oittrue**

This option is treated as synonymous with **-oi**.

**-oMa** <host address>

A number of options starting with **-oM** can be used to set values associated with remote hosts on locally-submitted messages (that is, messages not received over TCP/IP). These options can be used by any caller in conjunction with the **-bh**, **-be**, **-bf**, **-bF**, **-bt**, or **-bv** testing options. In other circumstances, they are ignored unless the caller is trusted.

The **-oMa** option sets the sender host address. This may include a port number at the end, after a full stop (period). For example:

```
exim -bs -oMa 10.9.8.7.1234
```

An alternative syntax is to enclose the IP address in square brackets, followed by a colon and the port number:

```
exim -bs -oMa [10.9.8.7]:1234
```

The IP address is placed in the *\$sender\_host\_address* variable, and the port, if present, in *\$sender\_host\_port*. If both **-oMa** and **-bh** are present on the command line, the sender host IP address is taken from whichever one is last.

**-oMaa** <name>

See **-oMa** above for general remarks about the **-oM** options. The **-oMaa** option sets the value of *\$sender\_host\_authenticated* (the authenticator name). See chapter 33 for a discussion of SMTP authentication. This option can be used with **-bh** and **-bs** to set up an authenticated SMTP session without actually using the SMTP AUTH command.

**-oMai** <string>

See **-oMa** above for general remarks about the **-oM** options. The **-oMai** option sets the value of *\$authenticated\_id* (the id that was authenticated). This overrides the default value (the caller's login id, except with **-bh**, where there is no default) for messages from local sources. See chapter 33 for a discussion of authenticated ids.

**-oMas** <address>

See **-oMa** above for general remarks about the **-oM** options. The **-oMas** option sets the authenticated sender value in *\$authenticated\_sender*. It overrides the sender address that is created from the caller's login id for messages from local sources, except when **-bh** is used, when there is no default. For both **-bh** and **-bs**, an authenticated sender that is specified on a MAIL command overrides this value. See chapter 33 for a discussion of authenticated senders.

**-oMi** <interface address>

See **-oMa** above for general remarks about the **-oM** options. The **-oMi** option sets the IP interface address value. A port number may be included, using the same syntax as for **-oMa**. The interface address is placed in *\$received\_ip\_address* and the port number, if present, in *\$received\_port*.

**-oMm** <message reference>

See **-oMa** above for general remarks about the **-oM** options. The **-oMm** option sets the message reference, e.g. message-id, and is logged during delivery. This is useful when some kind of audit trail is required to tie messages together. The format of the message reference is checked and will abort if the format is invalid. The option will only be accepted if exim is running in trusted mode, not as any regular user.

The best example of a message reference is when Exim sends a bounce message. The message reference is the message-id of the original message for which Exim is sending the bounce.

**-oMr** <protocol name>

See **-oMa** above for general remarks about the **-oM** options. The **-oMr** option sets the received protocol value that is stored in *\$received\_protocol*. However, it does not apply (and is ignored) when **-bh** or **-bs** is used. For **-bh**, the protocol is forced to one of the standard SMTP protocol names (see the description of *\$received\_protocol* in section 11.9). For **-bs**, the protocol is always “local-” followed by one of those same names. For **-bS** (batched SMTP) however, the protocol can be set by **-oMr**. Repeated use of this option is not supported.

**-oMs** <host name>

See **-oMa** above for general remarks about the **-oM** options. The **-oMs** option sets the sender host name in *\$sender\_host\_name*. When this option is present, Exim does not attempt to look up a host name from an IP address; it uses the name it is given.

**-oMt** <ident string>

See **-oMa** above for general remarks about the **-oM** options. The **-oMt** option sets the sender ident value in *\$sender\_ident*. The default setting for local callers is the login id of the calling process, except when **-bh** is used, when there is no default.

**-om**

In Sendmail, this option means “me too”, indicating that the sender of a message should receive a copy of the message if the sender appears in an alias expansion. Exim always does this, so the option does nothing.

**-oo**

This option is ignored. In Sendmail it specifies “old style headers”, whatever that means.

**-oP** <path>

This option is useful only in conjunction with **-bd** or **-q** with a time value. The option specifies the file to which the process id of the daemon is written. When **-oX** is used with **-bd**, or when **-q** with a time is used without **-bd**, this is the only way of causing Exim to write a pid file, because in those cases, the normal pid file is not used.

**-or** <time>

This option sets a timeout value for incoming non-SMTP messages. If it is not set, Exim will wait forever for the standard input. The value can also be set by the **receive\_timeout** option. The format used for specifying times is described in section 6.16.

**-os** <time>

This option sets a timeout value for incoming SMTP messages. The timeout applies to each SMTP command and block of data. The value can also be set by the **smtp\_receive\_timeout** option; it defaults to 5 minutes. The format used for specifying times is described in section 6.16.

**-ov**

This option has exactly the same effect as **-v**.

**-oX** <number or string>

This option is relevant only when the **-bd** (start listening daemon) option is also given. It controls which ports and interfaces the daemon uses. Details of the syntax, and how it interacts with configuration file options, are given in chapter 13. When **-oX** is used to start a daemon, no pid file is written unless **-oP** is also present to specify a pid file name.

## **-pd**

This option applies when an embedded Perl interpreter is linked with Exim (see chapter 12). It overrides the setting of the **perl\_at\_start** option, forcing the starting of the interpreter to be delayed until it is needed.

## **-ps**

This option applies when an embedded Perl interpreter is linked with Exim (see chapter 12). It overrides the setting of the **perl\_at\_start** option, forcing the starting of the interpreter to occur as soon as Exim is started.

## **-p<rval>:<sval>**

For compatibility with Sendmail, this option is equivalent to

```
-oMr <rval> -oMs <sval>
```

It sets the incoming protocol and host name (for trusted callers). The host name and its colon can be omitted when only the protocol is to be set. Note the Exim already has two private options, **-pd** and **-ps**, that refer to embedded Perl. It is therefore impossible to set a protocol value of **d** or **s** using this option (but that does not seem a real limitation). Repeated use of this option is not supported.

## **-q**

This option is normally restricted to admin users. However, there is a configuration option called **prod\_requires\_admin** which can be set false to relax this restriction (and also the same requirement for the **-M**, **-R**, and **-S** options).

If other commandline options do not specify an action, the **-q** option starts one queue runner process. This scans the queue of waiting messages, and runs a delivery process for each one in turn. It waits for each delivery process to finish before starting the next one. A delivery process may not actually do any deliveries if the retry times for the addresses have not been reached. Use **-qf** (see below) if you want to override this.

If the delivery process spawns other processes to deliver other messages down passed SMTP connections, the queue runner waits for these to finish before proceeding.

When all the queued messages have been considered, the original queue runner process terminates. In other words, a single pass is made over the waiting mail, one message at a time. Use **-q** with a time (see below) if you want this to be repeated periodically.

Exim processes the waiting messages in an unpredictable order. It isn't very random, but it is likely to be different each time, which is all that matters. If one particular message screws up a remote MTA, other messages to the same MTA have a chance of getting through if they get tried first.

It is possible to cause the messages to be processed in lexical message id order, which is essentially the order in which they arrived, by setting the **queue\_run\_in\_order** option, but this is not recommended for normal use.

## **-q<qflags>**

The **-q** option may be followed by one or more flag letters that change its behaviour. They are all optional, but if more than one is present, they must appear in the correct order. Each flag is described in a separate item below.

## **-qq...**

An option starting with **-qq** requests a two-stage queue run. In the first stage, the queue is scanned as if the **queue\_smtp\_domains** option matched every domain. Addresses are routed, local deliveries happen, but no remote transports are run.

The hints database that remembers which messages are waiting for specific hosts is updated, as if delivery to those hosts had been deferred. After this is complete, a second, normal queue scan happens, with routing and delivery taking place as normal. Messages that are routed to the same host should mostly be delivered down a single SMTP connection because of the hints that were set up during the first queue scan. This option may be useful for hosts that are connected to the Internet intermittently.

### **-q[q]i...**

If the *i* flag is present, the queue runner runs delivery processes only for those messages that haven't previously been tried. (*i* stands for "initial delivery".) This can be helpful if you are putting messages on the queue using **-odq** and want a queue runner just to process the new messages.

### **-q[q][i]f...**

If one *f* flag is present, a delivery attempt is forced for each non-frozen message, whereas without *f* only those non-frozen addresses that have passed their retry times are tried.

### **-q[q][i]ff...**

If *ff* is present, a delivery attempt is forced for every message, whether frozen or not.

### **-q[q][i][f]l**

The *l* (the letter "ell") flag specifies that only local deliveries are to be done. If a message requires any remote deliveries, it remains on the queue for later delivery.

### **-q[q][i][f]l[1][G<name>[/<time>]]**

If the *G* flag and a name is present, the queue runner operates on the queue with the given name rather than the default queue. The name should not contain a / character. For a periodic queue run (see below) append to the name a slash and a time value.

If other commandline options specify an action, a **-qG<name>** option will specify a queue to operate on. For example:

```
exim -bp -qGquarantine
mailq -qGquarantine
exim -qGoffpeak -Rf @special.domain.example
```

### **-q<qflags> <start id> <end id>**

When scanning the queue, Exim can be made to skip over messages whose ids are lexically less than a given value by following the **-q** option with a starting message id. For example:

```
exim -q 0t5C6f-0000c8-00
```

Messages that arrived earlier than 0t5C6f-0000c8-00 are not inspected. If a second message id is given, messages whose ids are lexically greater than it are also skipped. If the same id is given twice, for example,

```
exim -q 0t5C6f-0000c8-00 0t5C6f-0000c8-00
```

just one delivery process is started, for that message. This differs from **-M** in that retry data is respected, and it also differs from **-Mc** in that it counts as a delivery from a queue run. Note that the selection mechanism does not affect the order in which the messages are scanned. There are also other ways of selecting specific sets of messages for delivery in a queue run – see **-R** and **-S**.

### **-q<qflags><time>**

When a time value is present, the **-q** option causes Exim to run as a daemon, starting a queue runner process at intervals specified by the given time value (whose format is described in section 6.16). This form of the **-q** option is commonly combined with the **-bd** option, in which case a single daemon process handles both functions. A common way of starting up a combined daemon at system boot time is to use a command such as

```
/usr/exim/bin/exim -bd -q30m
```

Such a daemon listens for incoming SMTP calls, and also starts a queue runner process every 30 minutes.

When a daemon is started by **-q** with a time value, but without **-bd**, no pid file is written unless one is explicitly requested by the **-oP** option.

### **-qR<rsflags> <string>**

This option is synonymous with **-R**. It is provided for Sendmail compatibility.

### **-qS<rsflags> <string>**

This option is synonymous with **-S**.

### **-R**<rsflags> <string>

The <rsflags> may be empty, in which case the white space before the string is optional, unless the string is *f*, *ff*, *r*, *rf*, or *fff*, which are the possible values for <rsflags>. White space is required if <rsflags> is not empty.

This option is similar to **-q** with no time value, that is, it causes Exim to perform a single queue run, except that, when scanning the messages on the queue, Exim processes only those that have at least one undelivered recipient address containing the given string, which is checked in a case-independent way. If the <rsflags> start with *r*, <string> is interpreted as a regular expression; otherwise it is a literal string.

If you want to do periodic queue runs for messages with specific recipients, you can combine **-R** with **-q** and a time value. For example:

```
exim -q25m -R @special.domain.example
```

This example does a queue run for messages with recipients in the given domain every 25 minutes. Any additional flags that are specified with **-q** are applied to each queue run.

Once a message is selected for delivery by this mechanism, all its addresses are processed. For the first selected message, Exim overrides any retry information and forces a delivery attempt for each undelivered address. This means that if delivery of any address in the first message is successful, any existing retry information is deleted, and so delivery attempts for that address in subsequently selected messages (which are processed without forcing) will run. However, if delivery of any address does not succeed, the retry information is updated, and in subsequently selected messages, the failing address will be skipped.

If the <rsflags> contain *f* or *ff*, the delivery forcing applies to all selected messages, not just the first; frozen messages are included when *ff* is present.

The **-R** option makes it straightforward to initiate delivery of all messages to a given domain after a host has been down for some time. When the SMTP command ETRN is accepted by its ACL (see chapter 43), its default effect is to run Exim with the **-R** option, but it can be configured to run an arbitrary command instead.

### **-r**

This is a documented (for Sendmail) obsolete alternative name for **-f**.

### **-S**<rsflags> <string>

This option acts like **-R** except that it checks the string against each message's sender instead of against the recipients. If **-R** is also set, both conditions must be met for a message to be selected. If either of the options has *f* or *ff* in its flags, the associated action is taken.

### **-Tqt** <times>

This is an option that is exclusively for use by the Exim testing suite. It is not recognized when Exim is run normally. It allows for the setting up of explicit "queue times" so that various warning/retry features can be tested.

### **-t**

When Exim is receiving a locally-generated, non-SMTP message on its standard input, the **-t** option causes the recipients of the message to be obtained from the *To:*, *Cc:*, and *Bcc:* header lines in the message instead of from the command arguments. The addresses are extracted before any rewriting takes place and the *Bcc:* header line, if present, is then removed.

If the command has any arguments, they specify addresses to which the message is *not* to be delivered. That is, the argument addresses are removed from the recipients list obtained from the headers. This is compatible with Smail 3 and in accordance with the documented behaviour of several versions of Sendmail, as described in man pages on a number of operating systems (e.g. Solaris 8, IRIX 6.5, HP-UX 11). However, some versions of Sendmail *add* argument addresses to those obtained from the headers, and the O'Reilly Sendmail book documents it that way. Exim can be made to add argument addresses instead of subtracting them by setting the option **extract\_addresses\_remove\_arguments** false.

If there are any **Resent-** header lines in the message, Exim extracts recipients from all *Resent-To:*, *Resent-Cc:*, and *Resent-Bcc:* header lines instead of from *To:*, *Cc:*, and *Bcc:*. This is for compatibility with Sendmail and other MTAs. (Prior to release 4.20, Exim gave an error if **-t** was used in conjunction with **Resent-** header lines.)

RFC 2822 talks about different sets of **Resent-** header lines (for when a message is resent several times). The RFC also specifies that they should be added at the front of the message, and separated by *Received:* lines. It is not at all clear how **-t** should operate in the presence of multiple sets, nor indeed exactly what constitutes a “set”. In practice, it seems that MUAs do not follow the RFC. The **Resent-** lines are often added at the end of the header, and if a message is resent more than once, it is common for the original set of **Resent-** headers to be renamed as **X-Resent-** when a new set is added. This removes any possible ambiguity.

**-ti**

This option is exactly equivalent to **-t -i**. It is provided for compatibility with Sendmail.

**-tls-on-connect**

This option is available when Exim is compiled with TLS support. It forces all incoming SMTP connections to behave as if the incoming port is listed in the **tls\_on\_connect\_ports** option. See section 13.4 and chapter 42 for further details.

**-U**

Sendmail uses this option for “initial message submission”, and its documentation states that in future releases, it may complain about syntactically invalid messages rather than fixing them when this flag is not set. Exim ignores this option.

**-v**

This option causes Exim to write information to the standard error stream, describing what it is doing. In particular, it shows the log lines for receiving and delivering a message, and if an SMTP connection is made, the SMTP dialogue is shown. Some of the log lines shown may not actually be written to the log if the setting of **log\_selector** discards them. Any relevant selectors are shown with each log line. If none are shown, the logging is unconditional.

**-x**

AIX uses **-x** for a private purpose (“mail from a local mail program has National Language Support extended characters in the body of the mail item”). It sets **-x** when calling the MTA from its **mail** command. Exim ignores this option.

**-X <logfile>**

This option is interpreted by Sendmail to cause debug information to be sent to the named file. It is ignored by Exim.

**-z <log-line>**

This option writes its argument to Exim’s logfile. Use is restricted to administrators; the intent is for operational notes. Quotes should be used to maintain a multi-word item as a single argument, under most shells.



## 6. The Exim run time configuration file

Exim uses a single run time configuration file that is read whenever an Exim binary is executed. Note that in normal operation, this happens frequently, because Exim is designed to operate in a distributed manner, without central control.

If a syntax error is detected while reading the configuration file, Exim writes a message on the standard error, and exits with a non-zero return code. The message is also written to the panic log. **Note:** Only simple syntax errors can be detected at this time. The values of any expanded options are not checked until the expansion happens, even when the expansion does not actually alter the string.

The name of the configuration file is compiled into the binary for security reasons, and is specified by the `CONFIGURE_FILE` compilation option. In most configurations, this specifies a single file. However, it is permitted to give a colon-separated list of file names, in which case Exim uses the first existing file in the list.

The run time configuration file must be owned by root or by the user that is specified at compile time by the `CONFIGURE_OWNER` option (if set). The configuration file must not be world-writable, or group-writable unless its group is the root group or the one specified at compile time by the `CONFIGURE_GROUP` option.

**Warning:** In a conventional configuration, where the Exim binary is setuid to root, anybody who is able to edit the run time configuration file has an easy way to run commands as root. If you specify a user or group in the `CONFIGURE_OWNER` or `CONFIGURE_GROUP` options, then that user and/or any users who are members of that group will trivially be able to obtain root privileges.

Up to Exim version 4.72, the run time configuration file was also permitted to be writable by the Exim user and/or group. That has been changed in Exim 4.73 since it offered a simple privilege escalation for any attacker who managed to compromise the Exim user account.

A default configuration file, which will work correctly in simple situations, is provided in the file *src/configure.default*. If `CONFIGURE_FILE` defines just one file name, the installation process copies the default configuration to a new file of that name if it did not previously exist. If `CONFIGURE_FILE` is a list, no default is automatically installed. Chapter 7 is a “walk-through” discussion of the default configuration.

### 6.1 Using a different configuration file

A one-off alternate configuration can be specified by the `-C` command line option, which may specify a single file or a list of files. However, when `-C` is used, Exim gives up its root privilege, unless called by root (or unless the argument for `-C` is identical to the built-in value from `CONFIGURE_FILE`), or is listed in the `TRUSTED_CONFIG_LIST` file and the caller is the Exim user or the user specified in the `CONFIGURE_OWNER` setting. `-C` is useful mainly for checking the syntax of configuration files before installing them. No owner or group checks are done on a configuration file specified by `-C`, if root privilege has been dropped.

Even the Exim user is not trusted to specify an arbitrary configuration file with the `-C` option to be used with root privileges, unless that file is listed in the `TRUSTED_CONFIG_LIST` file. This locks out the possibility of testing a configuration using `-C` right through message reception and delivery, even if the caller is root. The reception works, but by that time, Exim is running as the Exim user, so when it re-execs to regain privilege for the delivery, the use of `-C` causes privilege to be lost. However, root can test reception and delivery using two separate commands (one to put a message on the queue, using `-odq`, and another to do the delivery, using `-M`).

If `ALT_CONFIG_PREFIX` is defined in *Local/Makefile*, it specifies a prefix string with which any file named in a `-C` command line option must start. In addition, the file name must not contain the sequence `“/ . . /”`. There is no default setting for `ALT_CONFIG_PREFIX`; when it is unset, any file name can be used with `-C`.

One-off changes to a configuration can be specified by the `-D` command line option, which defines and overrides values for macros used inside the configuration file. However, like `-C`, the use of this option by a non-privileged user causes Exim to discard its root privilege. If `DISABLE_D_OPTION` is

defined in *Local/Makefile*, the use of **-D** is completely disabled, and its use causes an immediate error exit.

The `WHITELIST_D_MACROS` option in *Local/Makefile* permits the binary builder to declare certain macro names trusted, such that root privilege will not necessarily be discarded. `WHITELIST_D_MACROS` defines a colon-separated list of macros which are considered safe and, if **-D** only supplies macros from this list, and the values are acceptable, then Exim will not give up root privilege if the caller is root, the Exim run-time user, or the `CONFIGURE_OWNER`, if set. This is a transition mechanism and is expected to be removed in the future. Acceptable values for the macros satisfy the regexp: `^[A-Za-z0-9_/.-]*$`

Some sites may wish to use the same Exim binary on different machines that share a file system, but to use different configuration files on each machine. If `CONFIGURE_FILE_USE_NODE` is defined in *Local/Makefile*, Exim first looks for a file whose name is the configuration file name followed by a dot and the machine's node name, as obtained from the `uname()` function. If this file does not exist, the standard name is tried. This processing occurs for each file name in the list given by `CONFIGURE_FILE` or **-C**.

In some esoteric situations different versions of Exim may be run under different effective uids and the `CONFIGURE_FILE_USE_EUID` is defined to help with this. See the comments in *src/EDITME* for details.

## 6.2 Configuration file format

Exim's configuration file is divided into a number of different parts. General option settings must always appear at the start of the file. The other parts are all optional, and may appear in any order. Each part other than the first is introduced by the word "begin" followed by at least one literal space, and the name of the part. The optional parts are:

- *ACL*: Access control lists for controlling incoming SMTP mail (see chapter 43).
- *authenticators*: Configuration settings for the authenticator drivers. These are concerned with the SMTP AUTH command (see chapter 33).
- *routers*: Configuration settings for the router drivers. Routers process addresses and determine how the message is to be delivered (see chapters 15–22).
- *transports*: Configuration settings for the transport drivers. Transports define mechanisms for copying messages to destinations (see chapters 24–30).
- *retry*: Retry rules, for use when a message cannot be delivered immediately. If there is no retry section, or if it is empty (that is, no retry rules are defined), Exim will not retry deliveries. In this situation, temporary errors are treated the same as permanent errors. Retry rules are discussed in chapter 32.
- *rewrite*: Global address rewriting rules, for use when a message arrives and when new addresses are generated during delivery. Rewriting is discussed in chapter 31.
- *local\_scan*: Private options for the `local_scan()` function. If you want to use this feature, you must set

```
LOCAL_SCAN_HAS_OPTIONS=yes
```

in *Local/Makefile* before building Exim. Details of the `local_scan()` facility are given in chapter 45.

Leading and trailing white space in configuration lines is always ignored.

Blank lines in the file, and lines starting with a `#` character (ignoring leading white space) are treated as comments and are ignored. **Note:** A `#` character other than at the beginning of a line is not treated specially, and does not introduce a comment.

Any non-comment line can be continued by ending it with a backslash. Note that the general rule for white space means that trailing white space after the backslash and leading white space at the start of continuation lines is ignored. Comment lines beginning with `#` (but not empty lines) may appear in the middle of a sequence of continuation lines.

A convenient way to create a configuration file is to start from the default, which is supplied in *src/configure.default*, and add, delete, or change settings as required.

The ACLs, retry rules, and rewriting rules have their own syntax which is described in chapters 43, 32, and 31, respectively. The other parts of the configuration file have some syntactic items in common, and these are described below, from section 6.11 onwards. Before that, the inclusion, macro, and conditional facilities are described.

## 6.3 File inclusions in the configuration file

You can include other files inside Exim's run time configuration file by using this syntax:

```
.include <file name>
.include_if_exists <file name>
```

on a line by itself. Double quotes round the file name are optional. If you use the first form, a configuration error occurs if the file does not exist; the second form does nothing for non-existent files. The first form allows a relative name. It is resolved relative to the directory of the including file. For the second form an absolute file name is required.

Includes may be nested to any depth, but remember that Exim reads its configuration file often, so it is a good idea to keep them to a minimum. If you change the contents of an included file, you must HUP the daemon, because an included file is read only when the configuration itself is read.

The processing of inclusions happens early, at a physical line level, so, like comment lines, an inclusion can be used in the middle of an option setting, for example:

```
hosts_lookup = a.b.c \
               .include /some/file
```

Include processing happens after macro processing (see below). Its effect is to process the lines of the included file as if they occurred inline where the inclusion appears.

## 6.4 Macros in the configuration file

If a line in the main part of the configuration (that is, before the first “begin” line) begins with an upper case letter, it is taken as a macro definition, and must be of the form

```
<name> = <rest of line>
```

The name must consist of letters, digits, and underscores, and need not all be in upper case, though that is recommended. The rest of the line, including any continuations, is the replacement text, and has leading and trailing white space removed. Quotes are not removed. The replacement text can never end with a backslash character, but this doesn't seem to be a serious limitation.

Macros may also be defined between router, transport, authenticator, or ACL definitions. They may not, however, be defined within an individual driver or ACL, or in the **local\_scan**, retry, or rewrite sections of the configuration.

## 6.5 Macro substitution

Once a macro is defined, all subsequent lines in the file (and any included files) are scanned for the macro name; if there are several macros, the line is scanned for each in turn, in the order in which the macros are defined. The replacement text is not re-scanned for the current macro, though it is scanned for subsequently defined macros. For this reason, a macro name may not contain the name of a previously defined macro as a substring. You could, for example, define

```
ABCD_XYZ = <something>
ABCD = <something else>
```

but putting the definitions in the opposite order would provoke a configuration error. Macro expansion is applied to individual physical lines from the file, before checking for line continuation or file inclusion (see above). If a line consists solely of a macro name, and the expansion of the macro is

empty, the line is ignored. A macro at the start of a line may turn the line into a comment line or a `.include` line.

## 6.6 Redefining macros

Once defined, the value of a macro can be redefined later in the configuration (or in an included file). Redefinition is specified by using `==` instead of `=`. For example:

```
MAC = initial value
...
MAC == updated value
```

Redefinition does not alter the order in which the macros are applied to the subsequent lines of the configuration file. It is still the same order in which the macros were originally defined. All that changes is the macro's value. Redefinition makes it possible to accumulate values. For example:

```
MAC = initial value
...
MAC == MAC and something added
```

This can be helpful in situations where the configuration file is built from a number of other files.

## 6.7 Overriding macro values

The values set for macros in the configuration file can be overridden by the **-D** command line option, but Exim gives up its root privilege when **-D** is used, unless called by root or the Exim user. A definition on the command line using the **-D** option causes all definitions and redefinitions within the file to be ignored.

## 6.8 Example of macro usage

As an example of macro usage, consider a configuration where aliases are looked up in a MySQL database. It helps to keep the file less cluttered if long strings such as SQL statements are defined separately as macros, for example:

```
ALIAS_QUERY = select mailbox from user where \
              login='${quote_mysql:$local_part}';
```

This can then be used in a *redirect* router setting like this:

```
data = ${lookup mysql{ALIAS_QUERY}}
```

In earlier versions of Exim macros were sometimes used for domain, host, or address lists. In Exim 4 these are handled better by named lists – see section 10.5.

## 6.9 Builtin macros

Exim defines some macros depending on facilities available, which may differ due to build-time definitions and from one release to another. All of these macros start with an underscore. They can be used to conditionally include parts of a configuration (see below).

The following classes of macros are defined:

<code>_HAVE_*</code>	build-time defines
<code>_DRIVER_ROUTER_*</code>	router drivers
<code>_DRIVER_TRANSPORT_*</code>	transport drivers
<code>_DRIVER_AUTHENTICATOR_*</code>	authenticator drivers
<code>_OPT_MAIN_*</code>	main config options
<code>_OPT_ROUTERS_*</code>	generic router options
<code>_OPT_TRANSPORTS_*</code>	generic transport options
<code>_OPT_AUTHENTICATORS_*</code>	generic authenticator options
<code>_OPT_ROUTER_*_*</code>	private router options

<code>_OPT_TRANSPORT_*_*</code>	private transport options
<code>_OPT_AUTHENTICATOR_*_*</code>	private authenticator options

Use an “`exim -bP macros`” command to get the list of macros.

## 6.10 Conditional skips in the configuration file

You can use the directives `.ifdef`, `.ifndef`, `.elifdef`, `.elifndef`, `.else`, and `.endif` to dynamically include or exclude portions of the configuration file. The processing happens whenever the file is read (that is, when an Exim binary starts to run).

The implementation is very simple. Instances of the first four directives must be followed by text that includes the names of one or macros. The condition that is tested is whether or not any macro substitution has taken place in the line. Thus:

```
.ifdef AAA
message_size_limit = 50M
.else
message_size_limit = 100M
.endif
```

sets a message size limit of 50M if the macro AAA is defined (or A or AA), and 100M otherwise. If there is more than one macro named on the line, the condition is true if any of them are defined. That is, it is an “or” condition. To obtain an “and” condition, you need to use nested `.ifdefs`.

Although you can use a macro expansion to generate one of these directives, it is not very useful, because the condition “there was a macro substitution in this line” will always be true.

Text following `.else` and `.endif` is ignored, and can be used as comment to clarify complicated nestings.

## 6.11 Common option syntax

For the main set of options, driver options, and *local\_scan()* options, each setting is on a line by itself, and starts with a name consisting of lower-case letters and underscores. Many options require a data value, and in these cases the name must be followed by an equals sign (with optional white space) and then the value. For example:

```
qualify_domain = mydomain.example.com
```

Some option settings may contain sensitive data, for example, passwords for accessing databases. To stop non-admin users from using the **-bP** command line option to read these values, you can precede the option settings with the word “hide”. For example:

```
hide mysql_servers = localhost/users/admin/secret-password
```

For non-admin users, such options are displayed like this:

```
mysql_servers = <value not displayable>
```

If “hide” is used on a driver option, it hides the value of that option on all instances of the same driver.

The following sections describe the syntax used for the different data types that are found in option settings.

## 6.12 Boolean options

Options whose type is given as boolean are on/off switches. There are two different ways of specifying such options: with and without a data value. If the option name is specified on its own without data, the switch is turned on; if it is preceded by “no\_” or “not\_” the switch is turned off. However, boolean options may be followed by an equals sign and one of the words “true”, “false”, “yes”, or “no”, as an alternative syntax. For example, the following two settings have exactly the same effect:

```
queue_only
queue_only = true
```

The following two lines also have the same (opposite) effect:

```
no_queue_only
queue_only = false
```

You can use whichever syntax you prefer.

## 6.13 Integer values

If an option's type is given as "integer", the value can be given in decimal, hexadecimal, or octal. If it starts with a digit greater than zero, a decimal number is assumed. Otherwise, it is treated as an octal number unless it starts with the characters "0x", in which case the remainder is interpreted as a hexadecimal number.

If an integer value is followed by the letter K, it is multiplied by 1024; if it is followed by the letter M, it is multiplied by 1024x1024; if by the letter G, 1024x1024x1024. When the values of integer option settings are output, values which are an exact multiple of 1024 or 1024x1024 are sometimes, but not always, printed using the letters K and M. The printing style is independent of the actual input format that was used.

## 6.14 Octal integer values

If an option's type is given as "octal integer", its value is always interpreted as an octal number, whether or not it starts with the digit zero. Such options are always output in octal.

## 6.15 Fixed point numbers

If an option's type is given as "fixed-point", its value must be a decimal integer, optionally followed by a decimal point and up to three further digits.

## 6.16 Time intervals

A time interval is specified as a sequence of numbers, each followed by one of the following letters, with no intervening white space:

<b>s</b>	seconds
<b>m</b>	minutes
<b>h</b>	hours
<b>d</b>	days
<b>w</b>	weeks

For example, "3h50m" specifies 3 hours and 50 minutes. The values of time intervals are output in the same format. Exim does not restrict the values; it is perfectly acceptable, for example, to specify "90m" instead of "1h30m".

## 6.17 String values

If an option's type is specified as "string", the value can be specified with or without double-quotes. If it does not start with a double-quote, the value consists of the remainder of the line plus any continuation lines, starting at the first character after any leading white space, with trailing white space removed, and with no interpretation of the characters in the string. Because Exim removes comment lines (those beginning with #) at an early stage, they can appear in the middle of a multi-line string. The following two settings are therefore equivalent:

```
trusted_users = uucp:mail
trusted_users = uucp:\
                  # This comment line is ignored
                  mail
```

If a string does start with a double-quote, it must end with a closing double-quote, and any backslash characters other than those used for line continuation are interpreted as escape characters, as follows:

<code>\\</code>	single backslash
<code>\n</code>	newline
<code>\r</code>	carriage return
<code>\t</code>	tab
<code>\&lt;octal digits&gt;</code>	up to 3 octal digits specify one character
<code>\x&lt;hex digits&gt;</code>	up to 2 hexadecimal digits specify one character

If a backslash is followed by some other character, including a double-quote character, that character replaces the pair.

Quoting is necessary only if you want to make use of the backslash escapes to insert special characters, or if you need to specify a value with leading or trailing spaces. These cases are rare, so quoting is almost never needed in current versions of Exim. In versions of Exim before 3.14, quoting was required in order to continue lines, so you may come across older configuration files and examples that apparently quote unnecessarily.

## 6.18 Expanded strings

Some strings in the configuration file are subjected to *string expansion*, by which means various parts of the string may be changed according to the circumstances (see chapter 11). The input syntax for such strings is as just described; in particular, the handling of backslashes in quoted strings is done as part of the input process, before expansion takes place. However, backslash is also an escape character for the expander, so any backslashes that are required for that reason must be doubled if they are within a quoted configuration string.

## 6.19 User and group names

User and group names are specified as strings, using the syntax described above, but the strings are interpreted specially. A user or group name must either consist entirely of digits, or be a name that can be looked up using the *getpwnam()* or *getgrnam()* function, as appropriate.

## 6.20 List construction

The data for some configuration options is a list of items, with colon as the default separator. Many of these options are shown with type “string list” in the descriptions later in this document. Others are listed as “domain list”, “host list”, “address list”, or “local part list”. Syntactically, they are all the same; however, those other than “string list” are subject to particular kinds of interpretation, as described in chapter 10.

In all these cases, the entire list is treated as a single string as far as the input syntax is concerned. The **trusted\_users** setting in section 6.17 above is an example. If a colon is actually needed in an item in a list, it must be entered as two colons. Leading and trailing white space on each item in a list is ignored. This makes it possible to include items that start with a colon, and in particular, certain forms of IPv6 address. For example, the list

```
local_interfaces = 127.0.0.1 : ::::1
```

contains two IP addresses, the IPv4 address 127.0.0.1 and the IPv6 address ::1.

**Note:** Although leading and trailing white space is ignored in individual list items, it is not ignored when parsing the list. The space after the first colon in the example above is necessary. If it were not there, the list would be interpreted as the two items 127.0.0.1:: and 1.

## 6.21 Changing list separators

Doubling colons in IPv6 addresses is an unwelcome chore, so a mechanism was introduced to allow the separator character to be changed. If a list begins with a left angle bracket, followed by any punctuation character, that character is used instead of colon as the list separator. For example, the list above can be rewritten to use a semicolon separator like this:

```
local_interfaces = <; 127.0.0.1 ; ::1
```

This facility applies to all lists, with the exception of the list in **log\_file\_path**. It is recommended that the use of non-colon separators be confined to circumstances where they really are needed.

It is also possible to use newline and other control characters (those with code values less than 32, plus DEL) as separators in lists. Such separators must be provided literally at the time the list is processed. For options that are string-expanded, you can write the separator using a normal escape sequence. This will be processed by the expander before the string is interpreted as a list. For example, if a newline-separated list of domains is generated by a lookup, you can process it directly by a line such as this:

```
domains = <\n ${lookup mysql{.....}}
```

This avoids having to change the list separator in such data. You are unlikely to want to use a control character as a separator in an option that is not expanded, because the value is literal text. However, it can be done by giving the value in quotes. For example:

```
local_interfaces = "<\n 127.0.0.1 \n ::1"
```

Unlike printing character separators, which can be included in list items by doubling, it is not possible to include a control character as data when it is set as the separator. Two such characters in succession are interpreted as enclosing an empty list item.

## 6.22 Empty items in lists

An empty item at the end of a list is always ignored. In other words, trailing separator characters are ignored. Thus, the list in

```
senders = user@domain :
```

contains only a single item. If you want to include an empty string as one item in a list, it must not be the last item. For example, this list contains three items, the second of which is empty:

```
senders = user1@domain : : user2@domain
```

**Note:** There must be white space between the two colons, as otherwise they are interpreted as representing a single colon data character (and the list would then contain just one item). If you want to specify a list that contains just one, empty item, you can do it as in this example:

```
senders = :
```

In this case, the first item is empty, and the second is discarded because it is at the end of the list.

## 6.23 Format of driver configurations

There are separate parts in the configuration for defining routers, transports, and authenticators. In each part, you are defining a number of driver instances, each with its own set of options. Each driver instance is defined by a sequence of lines like this:

```
<instance name>:  
<option>  
...  
<option>
```

In the following example, the instance name is *localuser*, and it is followed by three options settings:

```
localuser:  
  driver = accept  
  check_local_user  
  transport = local_delivery
```

For each driver instance, you specify which Exim code module it uses – by the setting of the **driver** option – and (optionally) some configuration settings. For example, in the case of transports, if you want a transport to deliver with SMTP you would use the *smtp* driver; if you want to deliver to a local file you would use the *appendfile* driver. Each of the drivers is described in detail in its own separate chapter later in this manual.



You can have several routers, transports, or authenticators that are based on the same underlying driver (each must have a different instance name).

The order in which routers are defined is important, because addresses are passed to individual routers one by one, in order. The order in which transports are defined does not matter at all. The order in which authenticators are defined is used only when Exim, as a client, is searching them to find one that matches an authentication mechanism offered by the server.

Within a driver instance definition, there are two kinds of option: *generic* and *private*. The generic options are those that apply to all drivers of the same type (that is, all routers, all transports or all authenticators). The **driver** option is a generic option that must appear in every definition. The private options are special for each driver, and none need appear, because they all have default values.

The options may appear in any order, except that the **driver** option must precede any private options, since these depend on the particular driver. For this reason, it is recommended that **driver** always be the first option.

Driver instance names, which are used for reference in log entries and elsewhere, can be any sequence of letters, digits, and underscores (starting with a letter) and must be unique among drivers of the same type. A router and a transport (for example) can each have the same name, but no two router instances can have the same name. The name of a driver instance should not be confused with the name of the underlying driver module. For example, the configuration lines:

```
remote_smtp:
    driver = smtp
```

create an instance of the *smtp* transport driver whose name is *remote\_smtp*. The same driver code can be used more than once, with different instance names and different option settings each time. A second instance of the *smtp* transport, with different options, might be defined thus:

```
special_smtp:
    driver = smtp
    port = 1234
    command_timeout = 10s
```

The names *remote\_smtp* and *special\_smtp* would be used to reference these transport instances from routers, and these names would appear in log lines.

Comment lines may be present in the middle of driver specifications. The full list of option settings for any particular driver instance, including all the defaulted values, can be extracted by making use of the **-bP** command line option.

## 7. The default configuration file

The default configuration file supplied with Exim as *src/configure.default* is sufficient for a host with simple mail requirements. As an introduction to the way Exim is configured, this chapter “walks through” the default configuration, giving brief explanations of the settings. Detailed descriptions of the options are given in subsequent chapters. The default configuration file itself contains extensive comments about ways you might want to modify the initial settings. However, note that there are many options that are not mentioned at all in the default configuration.

### 7.1 Main configuration settings

The main (global) configuration option settings must always come first in the file. The first thing you’ll see in the file, after some initial comments, is the line

```
# primary_hostname =
```

This is a commented-out setting of the **primary\_hostname** option. Exim needs to know the official, fully qualified name of your host, and this is where you can specify it. However, in most cases you do not need to set this option. When it is unset, Exim uses the *uname()* system function to obtain the host name.

The first three non-comment configuration lines are as follows:

```
domainlist local_domains      = @
domainlist relay_to_domains   =
hostlist   relay_from_hosts   = 127.0.0.1
```

These are not, in fact, option settings. They are definitions of two named domain lists and one named host list. Exim allows you to give names to lists of domains, hosts, and email addresses, in order to make it easier to manage the configuration file (see section 10.5).

The first line defines a domain list called *local\_domains*; this is used later in the configuration to identify domains that are to be delivered on the local host.

There is just one item in this list, the string “@”. This is a special form of entry which means “the name of the local host”. Thus, if the local host is called *a.host.example*, mail to *any.user@a.host.example* is expected to be delivered locally. Because the local host’s name is referenced indirectly, the same configuration file can be used on different hosts.

The second line defines a domain list called *relay\_to\_domains*, but the list itself is empty. Later in the configuration we will come to the part that controls mail relaying through the local host; it allows relaying to any domains in this list. By default, therefore, no relaying on the basis of a mail domain is permitted.

The third line defines a host list called *relay\_from\_hosts*. This list is used later in the configuration to permit relaying from any host or IP address that matches the list. The default contains just the IP address of the IPv4 loopback interface, which means that processes on the local host are able to submit mail for relaying by sending it over TCP/IP to that interface. No other hosts are permitted to submit messages for relaying.

Just to be sure there’s no misunderstanding: at this point in the configuration we aren’t actually setting up any controls. We are just defining some domains and hosts that will be used in the controls that are specified later.

The next two configuration lines are genuine option settings:

```
acl_smtp_rcpt = acl_check_rcpt
acl_smtp_data = acl_check_data
```

These options specify *Access Control Lists* (ACLs) that are to be used during an incoming SMTP session for every recipient of a message (every RCPT command), and after the contents of the message have been received, respectively. The names of the lists are *acl\_check\_rcpt* and *acl\_check\_data*, and we will come to their definitions below, in the ACL section of the configuration. The RCPT ACL controls which recipients are accepted for an incoming message – if a configuration

does not provide an ACL to check recipients, no SMTP mail can be accepted. The DATA ACL allows the contents of a message to be checked.

Two commented-out option settings are next:

```
# av_scanner = clamd:/tmp/clamd
# spamd_address = 127.0.0.1 783
```

These are example settings that can be used when Exim is compiled with the content-scanning extension. The first specifies the interface to the virus scanner, and the second specifies the interface to SpamAssassin. Further details are given in chapter 44.

Three more commented-out option settings follow:

```
# tls_advertise_hosts = *
# tls_certificate = /etc/ssl/exim.crt
# tls_privatekey = /etc/ssl/exim.pem
```

These are example settings that can be used when Exim is compiled with support for TLS (aka SSL) as described in section 4.7. The first one specifies the list of clients that are allowed to use TLS when connecting to this server; in this case the wildcard means all clients. The other options specify where Exim should find its TLS certificate and private key, which together prove the server's identity to any clients that connect. More details are given in chapter 42.

Another two commented-out option settings follow:

```
# daemon_smtp_ports = 25 : 465 : 587
# tls_on_connect_ports = 465
```

These options provide better support for roaming users who wish to use this server for message submission. They are not much use unless you have turned on TLS (as described in the previous paragraph) and authentication (about which more in section 7.7). The usual SMTP port 25 is often blocked on end-user networks, so RFC 4409 specifies that message submission should use port 587 instead. However some software (notably Microsoft Outlook) cannot be configured to use port 587 correctly, so these settings also enable the non-standard “smtps” (aka “ssmtp”) port 465 (see section 13.4).

Two more commented-out options settings follow:

```
# qualify_domain =
# qualify_recipient =
```

The first of these specifies a domain that Exim uses when it constructs a complete email address from a local login name. This is often needed when Exim receives a message from a local process. If you do not set **qualify\_domain**, the value of **primary\_hostname** is used. If you set both of these options, you can have different qualification domains for sender and recipient addresses. If you set only the first one, its value is used in both cases.

The following line must be uncommented if you want Exim to recognize addresses of the form *user@[10.11.12.13]* that is, with a “domain literal” (an IP address within square brackets) instead of a named domain.

```
# allow_domain_literals
```

The RFCs still require this form, but many people think that in the modern Internet it makes little sense to permit mail to be sent to specific hosts by quoting their IP addresses. This ancient format has been used by people who try to abuse hosts by using them for unwanted relaying. However, some people believe there are circumstances (for example, messages addressed to *postmaster*) where domain literals are still useful.

The next configuration line is a kind of trigger guard:

```
never_users = root
```

It specifies that no delivery must ever be run as the root user. The normal convention is to set up *root* as an alias for the system administrator. This setting is a guard against slips in the configuration. The list of users specified by **never\_users** is not, however, the complete list; the build-time configuration

in *Local/Makefile* has an option called `FIXED_NEVER_USERS` specifying a list that cannot be overridden. The contents of `never_users` are added to this list. By default `FIXED_NEVER_USERS` also specifies root.

When a remote host connects to Exim in order to send mail, the only information Exim has about the host's identity is its IP address. The next configuration line,

```
host_lookup = *
```

specifies that Exim should do a reverse DNS lookup on all incoming connections, in order to get a host name. This improves the quality of the logging information, but if you feel it is too expensive, you can remove it entirely, or restrict the lookup to hosts on “nearby” networks. Note that it is not always possible to find a host name from an IP address, because not all DNS reverse zones are maintained, and sometimes DNS servers are unreachable.

The next two lines are concerned with *ident* callbacks, as defined by RFC 1413 (hence their names):

```
rfc1413_hosts = *
rfc1413_query_timeout = 0s
```

These settings cause Exim to avoid ident callbacks for all incoming SMTP calls. Few hosts offer RFC1413 service these days; calls have to be terminated by a timeout and this needlessly delays the startup of an incoming SMTP connection. If you have hosts for which you trust RFC1413 and need this information, you can change this.

This line enables an efficiency SMTP option. It is negotiated by clients and not expected to cause problems but can be disabled if needed.

```
prdr_enable = true
```

When Exim receives messages over SMTP connections, it expects all addresses to be fully qualified with a domain, as required by the SMTP definition. However, if you are running a server to which simple clients submit messages, you may find that they send unqualified addresses. The two commented-out options:

```
# sender_unqualified_hosts =
# recipient_unqualified_hosts =
```

show how you can specify hosts that are permitted to send unqualified sender and recipient addresses, respectively.

The **log\_selector** option is used to increase the detail of logging over the default:

```
log_selector = +smtp_protocol_error +smtp_syntax_error \
               +tls_certificate_verified
```

The **percent\_hack\_domains** option is also commented out:

```
# percent_hack_domains =
```

It provides a list of domains for which the “percent hack” is to operate. This is an almost obsolete form of explicit email routing. If you do not know anything about it, you can safely ignore this topic.

The next two settings in the main part of the default configuration are concerned with messages that have been “frozen” on Exim's queue. When a message is frozen, Exim no longer continues to try to deliver it. Freezing occurs when a bounce message encounters a permanent failure because the sender address of the original message that caused the bounce is invalid, so the bounce cannot be delivered. This is probably the most common case, but there are also other conditions that cause freezing, and frozen messages are not always bounce messages.

```
ignore_bounce_errors_after = 2d
timeout_frozen_after = 7d
```

The first of these options specifies that failing bounce messages are to be discarded after 2 days on the queue. The second specifies that any frozen message (whether a bounce message or not) is to be timed out (and discarded) after a week. In this configuration, the first setting ensures that no failing bounce message ever lasts a week.

Exim queues its messages in a spool directory. If you expect to have large queues, you may consider using this option. It splits the spool directory into subdirectories to avoid file system degradation from many files in a single directory, resulting in better performance. Manual manipulation of queued messages becomes more complex (though fortunately not often needed).

```
# split_spool_directory = true
```

In an ideal world everybody follows the standards. For non-ASCII messages RFC 2047 is a standard, allowing a maximum line length of 76 characters. Exim adheres that standard and won't process messages which violate this standard. (Even `{rfc2047:...}` expansions will fail.) In particular, the Exim maintainers have had multiple reports of problems from Russian administrators of issues until they disable this check, because of some popular, yet buggy, mail composition software.

```
# check_rfc2047_length = false
```

If you need to be strictly RFC compliant you may wish to disable the 8BITMIME advertisement. Use this, if you exchange mails with systems that are not 8-bit clean.

```
# accept_8bitmime = false
```

Libraries you use may depend on specific environment settings. This imposes a security risk (e.g. `PATH`). There are two lists: **keep\_environment** for the variables to import as they are, and **add\_environment** for variables we want to set to a fixed value. Note that `TZ` is handled separately, by the `$(timezone%)` runtime option and by the `TIMEZONE_DEFAULT` buildtime option.

```
# keep_environment = ^LDAP
# add_environment = PATH=/usr/bin:./bin
```

## 7.2 ACL configuration

In the default configuration, the ACL section follows the main configuration. It starts with the line

```
begin acl
```

and it contains the definitions of two ACLs, called *acl\_check\_rcpt* and *acl\_check\_data*, that were referenced in the settings of **acl\_smtp\_rcpt** and **acl\_smtp\_data** above.

The first ACL is used for every RCPT command in an incoming SMTP message. Each RCPT command specifies one of the message's recipients. The ACL statements are considered in order, until the recipient address is either accepted or rejected. The RCPT command is then accepted or rejected, according to the result of the ACL processing.

```
acl_check_rcpt:
```

This line, consisting of a name terminated by a colon, marks the start of the ACL, and names it.

```
accept hosts = :
```

This ACL statement accepts the recipient if the sending host matches the list. But what does that strange list mean? It doesn't actually contain any host names or IP addresses. The presence of the colon puts an empty item in the list; Exim matches this only if the incoming message did not come from a remote host, because in that case, the remote hostname is empty. The colon is important. Without it, the list itself is empty, and can never match anything.

What this statement is doing is to accept unconditionally all recipients in messages that are submitted by SMTP from local processes using the standard input and output (that is, not using TCP/IP). A number of MUAs operate in this manner.

```
deny      message      = Restricted characters in address
          domains      = +local_domains
          local_parts   = ^[.] : ^.*[@%!/|]

deny      message      = Restricted characters in address
          domains      = !+local_domains
          local_parts   = ^[./|] : ^.*[@%!|] : ^.*\\/\\.\\.\\.\/
```

These statements are concerned with local parts that contain any of the characters “@”, “%”, “!”, “/”, “[”, or dots in unusual places. Although these characters are entirely legal in local parts (in the case of “@” and leading dots, only if correctly quoted), they do not commonly occur in Internet mail addresses.

The first three have in the past been associated with explicitly routed addresses (percent is still sometimes used – see the **percent\_hack\_domains** option). Addresses containing these characters are regularly tried by spammers in an attempt to bypass relaying restrictions, and also by open relay testing programs. Unless you really need them it is safest to reject these characters at this early stage. This configuration is heavy-handed in rejecting these characters for all messages it accepts from remote hosts. This is a deliberate policy of being as safe as possible.

The first rule above is stricter, and is applied to messages that are addressed to one of the local domains handled by this host. This is implemented by the first condition, which restricts it to domains that are listed in the *local\_domains* domain list. The “+” character is used to indicate a reference to a named list. In this configuration, there is just one domain in *local\_domains*, but in general there may be many.

The second condition on the first statement uses two regular expressions to block local parts that begin with a dot or contain “@”, “%”, “!”, “/”, or “[”. If you have local accounts that include these characters, you will have to modify this rule.

Empty components (two dots in a row) are not valid in RFC 2822, but Exim allows them because they have been encountered in practice. (Consider the common convention of local parts constructed as “*first-initial.second-initial.family-name*” when applied to someone like the author of Exim, who has no second initial.) However, a local part starting with a dot or containing “*/. /*” can cause trouble if it is used as part of a file name (for example, for a mailing list). This is also true for local parts that contain slashes. A pipe symbol can also be troublesome if the local part is incorporated unthinkingly into a shell command line.

The second rule above applies to all other domains, and is less strict. This allows your own users to send outgoing messages to sites that use slashes and vertical bars in their local parts. It blocks local parts that begin with a dot, slash, or vertical bar, but allows these characters within the local part. However, the sequence “*/. /*” is barred. The use of “@”, “%”, and “!” is blocked, as before. The motivation here is to prevent your users (or your users’ viruses) from mounting certain kinds of attack on remote sites.

```
accept local_parts    = postmaster
       domains        = +local_domains
```

This statement, which has two conditions, accepts an incoming address if the local part is *postmaster* and the domain is one of those listed in the *local\_domains* domain list. The “+” character is used to indicate a reference to a named list. In this configuration, there is just one domain in *local\_domains*, but in general there may be many.

The presence of this statement means that mail to *postmaster* is never blocked by any of the subsequent tests. This can be helpful while sorting out problems in cases where the subsequent tests are incorrectly denying access.

```
require verify        = sender
```

This statement requires the sender address to be verified before any subsequent ACL statement can be used. If verification fails, the incoming recipient address is refused. Verification consists of trying to route the address, to see if a bounce message could be delivered to it. In the case of remote addresses, basic verification checks only the domain, but *callouts* can be used for more verification if required. Section 43.44 discusses the details of address verification.

```
accept hosts          = +relay_from_hosts
       control         = submission
```

This statement accepts the address if the message is coming from one of the hosts that are defined as being allowed to relay through this host. Recipient verification is omitted here, because in many cases the clients are dumb MUAs that do not cope well with SMTP error responses. For the same reason, the second line specifies “*submission mode*” for messages that are accepted. This is described in

detail in section 47.1; it causes Exim to fix messages that are deficient in some way, for example, because they lack a *Date:* header line. If you are actually relaying out from MTAs, you should probably add recipient verification here, and disable submission mode.

```
accept authenticated = *
control              = submission
```

This statement accepts the address if the client host has authenticated itself. Submission mode is again specified, on the grounds that such messages are most likely to come from MUAs. The default configuration does not define any authenticators, though it does include some nearly complete commented-out examples described in 7.7. This means that no client can in fact authenticate until you complete the authenticator definitions.

```
require message = relay not permitted
domains = +local_domains : +relay_to_domains
```

This statement rejects the address if its domain is neither a local domain nor one of the domains for which this host is a relay.

```
require verify = recipient
```

This statement requires the recipient address to be verified; if verification fails, the address is rejected.

```
# deny      message      = rejected because $sender_host_address \
#                                     is in a black list at $dnslist_domain\n\
#                                     $dnslist_text
#          dnslists      = black.list.example
#
# warn      dnslists      = black.list.example
#          add_header     = X-Warning: $sender_host_address is in \
#                                     a black list at $dnslist_domain
#          log_message    = found in $dnslist_domain
```

These commented-out lines are examples of how you could configure Exim to check sending hosts against a DNS black list. The first statement rejects messages from blacklisted hosts, whereas the second just inserts a warning header line.

```
# require verify = csa
```

This commented-out line is an example of how you could turn on client SMTP authorization (CSA) checking. Such checks do DNS lookups for special SRV records.

```
accept
```

The final statement in the first ACL unconditionally accepts any recipient address that has successfully passed all the previous tests.

```
acl_check_data:
```

This line marks the start of the second ACL, and names it. Most of the contents of this ACL are commented out:

```
# deny      malware      = *
#          message      = This message contains a virus \
#                                     ($malware_name).
```

These lines are examples of how to arrange for messages to be scanned for viruses when Exim has been compiled with the content-scanning extension, and a suitable virus scanner is installed. If the message is found to contain a virus, it is rejected with the given custom error message.

```
# warn      spam          = nobody
#          message      = X-Spam_score: $spam_score\n\
#                                     X-Spam_score_int: $spam_score_int\n\
#                                     X-Spam_bar: $spam_bar\n\
#                                     X-Spam_report: $spam_report
```

These lines are an example of how to arrange for messages to be scanned by SpamAssassin when Exim has been compiled with the content-scanning extension, and SpamAssassin has been installed. The SpamAssassin check is run with `nobody` as its user parameter, and the results are added to the message as a series of extra header line. In this case, the message is not rejected, whatever the spam score.

```
accept
```

This final line in the DATA ACL accepts the message unconditionally.

## 7.3 Router configuration

The router configuration comes next in the default configuration, introduced by the line

```
begin routers
```

Routers are the modules in Exim that make decisions about where to send messages. An address is passed to each router in turn, until it is either accepted, or failed. This means that the order in which you define the routers matters. Each router is fully described in its own chapter later in this manual. Here we give only brief overviews.

```
# domain_literal:
#   driver = ipliteral
#   domains = !+local_domains
#   transport = remote_smtp
```

This router is commented out because the majority of sites do not want to support domain literal addresses (those of the form `user@[10.9.8.7]`). If you uncomment this router, you also need to uncomment the setting of **allow\_domain\_literals** in the main part of the configuration.

```
dnslookup:
  driver = dnslookup
  domains = ! +local_domains
  transport = remote_smtp
  ignore_target_hosts = 0.0.0.0 : 127.0.0.0/8
  no_more
```

The first uncommented router handles addresses that do not involve any local domains. This is specified by the line

```
domains = ! +local_domains
```

The **domains** option lists the domains to which this router applies, but the exclamation mark is a negation sign, so the router is used only for domains that are not in the domain list called *local\_domains* (which was defined at the start of the configuration). The plus sign before *local\_domains* indicates that it is referring to a named list. Addresses in other domains are passed on to the following routers.

The name of the router driver is *dnslookup*, and is specified by the **driver** option. Do not be confused by the fact that the name of this router instance is the same as the name of the driver. The instance name is arbitrary, but the name set in the **driver** option must be one of the driver modules that is in the Exim binary.

The *dnslookup* router routes addresses by looking up their domains in the DNS in order to obtain a list of hosts to which the address is routed. If the router succeeds, the address is queued for the *remote\_smtp* transport, as specified by the **transport** option. If the router does not find the domain in the DNS, no further routers are tried because of the **no\_more** setting, so the address fails and is bounced.

The **ignore\_target\_hosts** option specifies a list of IP addresses that are to be entirely ignored. This option is present because a number of cases have been encountered where MX records in the DNS point to host names whose IP addresses are 0.0.0.0 or are in the 127 subnet (typically 127.0.0.1). Completely ignoring these IP addresses causes Exim to fail to route the email address, so it bounces.



Otherwise, Exim would log a routing problem, and continue to try to deliver the message periodically until the address timed out.

```
system_aliases:
  driver = redirect
  allow_fail
  allow_defer
  data = ${lookup{$local_part}lsearch{/etc/aliases}}
# user = exim
  file_transport = address_file
  pipe_transport = address_pipe
```

Control reaches this and subsequent routers only for addresses in the local domains. This router checks to see whether the local part is defined as an alias in the */etc/aliases* file, and if so, redirects it according to the data that it looks up from that file. If no data is found for the local part, the value of the **data** option is empty, causing the address to be passed to the next router.

*/etc/aliases* is a conventional name for the system aliases file that is often used. That is why it is referenced by from the default configuration file. However, you can change this by setting **SYSTEM\_ALIASES\_FILE** in *Local/Makefile* before building Exim.

```
userforward:
  driver = redirect
  check_local_user
# local_part_suffix = +* : -*
# local_part_suffix_optional
  file = $home/.forward
# allow_filter
  no_verify
  no_expn
  check_ancestor
  file_transport = address_file
  pipe_transport = address_pipe
  reply_transport = address_reply
```

This is the most complicated router in the default configuration. It is another redirection router, but this time it is looking for forwarding data set up by individual users. The **check\_local\_user** setting specifies a check that the local part of the address is the login name of a local user. If it is not, the router is skipped. The two commented options that follow **check\_local\_user**, namely:

```
# local_part_suffix = +* : -*
# local_part_suffix_optional
```

show how you can specify the recognition of local part suffixes. If the first is uncommented, a suffix beginning with either a plus or a minus sign, followed by any sequence of characters, is removed from the local part and placed in the variable *\$local\_part\_suffix*. The second suffix option specifies that the presence of a suffix in the local part is optional. When a suffix is present, the check for a local login uses the local part with the suffix removed.

When a local user account is found, the file called *.forward* in the user's home directory is consulted. If it does not exist, or is empty, the router declines. Otherwise, the contents of *.forward* are interpreted as redirection data (see chapter 22 for more details).

Traditional *.forward* files contain just a list of addresses, pipes, or files. Exim supports this by default. However, if **allow\_filter** is set (it is commented out by default), the contents of the file are interpreted as a set of Exim or Sieve filtering instructions, provided the file begins with “#Exim filter” or “#Sieve filter”, respectively. User filtering is discussed in the separate document entitled *Exim's interfaces to mail filtering*.

The **no\_verify** and **no\_expn** options mean that this router is skipped when verifying addresses, or when running as a consequence of an SMTP EXPN command. There are two reasons for doing this:

- (1) Whether or not a local user has a *.forward* file is not really relevant when checking an address for validity; it makes sense not to waste resources doing unnecessary work.
- (2) More importantly, when Exim is verifying addresses or handling an EXPN command during an SMTP session, it is running as the Exim user, not as root. The group is the Exim group, and no additional groups are set up. It may therefore not be possible for Exim to read users' *.forward* files at this time.

The setting of **check\_ancestor** prevents the router from generating a new address that is the same as any previous address that was redirected. (This works round a problem concerning a bad interaction between aliasing and forwarding – see section 22.5).

The final three option settings specify the transports that are to be used when forwarding generates a direct delivery to a file, or to a pipe, or sets up an auto-reply, respectively. For example, if a *.forward* file contains

```
a.nother@elsewhere.example, /home/spqr/archive
```

the delivery to */home/spqr/archive* is done by running the **address\_file** transport.

```
localuser:
  driver = accept
  check_local_user
# local_part_suffix = +* : -*
# local_part_suffix_optional
  transport = local_delivery
```

The final router sets up delivery into local mailboxes, provided that the local part is the name of a local login, by accepting the address and assigning it to the *local\_delivery* transport. Otherwise, we have reached the end of the routers, so the address is bounced. The commented suffix settings fulfil the same purpose as they do for the *userforward* router.

## 7.4 Transport configuration

Transports define mechanisms for actually delivering messages. They operate only when referenced from routers, so the order in which they are defined does not matter. The transports section of the configuration starts with

```
begin transports
```

One remote transport and four local transports are defined.

```
remote_smtp:
  driver = smtp
  hosts_try_prdr = *
```

This transport is used for delivering messages over SMTP connections. The list of remote hosts comes from the router. The **hosts\_try\_prdr** option enables an efficiency SMTP option. It is negotiated between client and server and not expected to cause problems but can be disabled if needed. All other options are defaulted.

```
local_delivery:
  driver = appendfile
  file = /var/mail/$local_part
  delivery_date_add
  envelope_to_add
  return_path_add
# group = mail
# mode = 0660
```

This *appendfile* transport is used for local delivery to user mailboxes in traditional BSD mailbox format. By default it runs under the uid and gid of the local user, which requires the sticky bit to be set on the */var/mail* directory. Some systems use the alternative approach of running mail deliveries under a particular group instead of using the sticky bit. The commented options show how this can be done.

Exim adds three headers to the message as it delivers it: *Delivery-date:*, *Envelope-to:* and *Return-path:*. This action is requested by the three similarly-named options above.

```
address_pipe:
    driver = pipe
    return_output
```

This transport is used for handling deliveries to pipes that are generated by redirection (aliasing or users' *.forward* files). The **return\_output** option specifies that any output on stdout or stderr generated by the pipe is to be returned to the sender.

```
address_file:
    driver = appendfile
    delivery_date_add
    envelope_to_add
    return_path_add
```

This transport is used for handling deliveries to files that are generated by redirection. The name of the file is not specified in this instance of *appendfile*, because it comes from the *redirect* router.

```
address_reply:
    driver = autoreply
```

This transport is used for handling automatic replies generated by users' filter files.

## 7.5 Default retry rule

The retry section of the configuration file contains rules which affect the way Exim retries deliveries that cannot be completed at the first attempt. It is introduced by the line

```
begin retry
```

In the default configuration, there is just one rule, which applies to all errors:

```
*      *      F,2h,15m; G,16h,1h,1.5; F,4d,6h
```

This causes any temporarily failing address to be retried every 15 minutes for 2 hours, then at intervals starting at one hour and increasing by a factor of 1.5 until 16 hours have passed, then every 6 hours up to 4 days. If an address is not delivered after 4 days of temporary failure, it is bounced. The time is measured from first failure, not from the time the message was received.

If the retry section is removed from the configuration, or is empty (that is, if no retry rules are defined), Exim will not retry deliveries. This turns temporary errors into permanent errors.

## 7.6 Rewriting configuration

The rewriting section of the configuration, introduced by

```
begin rewrite
```

contains rules for rewriting addresses in messages as they arrive. There are no rewriting rules in the default configuration file.

## 7.7 Authenticators configuration

The authenticators section of the configuration, introduced by

```
begin authenticators
```

defines mechanisms for the use of the SMTP AUTH command. The default configuration file contains two commented-out example authenticators which support plaintext username/password authentication using the standard PLAIN mechanism and the traditional but non-standard LOGIN mechanism, with Exim acting as the server. PLAIN and LOGIN are enough to support most MUA software.

The example PLAIN authenticator looks like this:

```
#PLAIN:
#  driver                      = plaintext
#  server_set_id               = $auth2
#  server_prompts              = :
#  server_condition            = Authentication is not yet configured
#  server_advertise_condition = ${if def:tls_in_cipher }
```

And the example LOGIN authenticator looks like this:

```
#LOGIN:
#  driver                      = plaintext
#  server_set_id               = $auth1
#  server_prompts              = <| Username: | Password:
#  server_condition            = Authentication is not yet configured
#  server_advertise_condition = ${if def:tls_in_cipher }
```

The **server\_set\_id** option makes Exim remember the authenticated username in *\$authenticated\_id*, which can be used later in ACLs or routers. The **server\_prompts** option configures the *plaintext* authenticator so that it implements the details of the specific authentication mechanism, i.e. PLAIN or LOGIN. The **server\_advertise\_condition** setting controls when Exim offers authentication to clients; in the examples, this is only when TLS or SSL has been started, so to enable the authenticators you also need to add support for TLS as described in section 7.1.

The **server\_condition** setting defines how to verify that the username and password are correct. In the examples it just produces an error message. To make the authenticators work, you can use a string expansion expression like one of the examples in chapter 34.

Beware that the sequence of the parameters to PLAIN and LOGIN differ; the usercode and password are in different positions. Chapter 34 covers both.

## 8. Regular expressions

Exim supports the use of regular expressions in many of its options. It uses the PCRE regular expression library; this provides regular expression matching that is compatible with Perl 5. The syntax and semantics of regular expressions is discussed in online Perl manpages, in many Perl reference books, and also in Jeffrey Friedl's *Mastering Regular Expressions*, which is published by O'Reilly (see <http://www.oreilly.com/catalog/regex2/>).

The documentation for the syntax and semantics of the regular expressions that are supported by PCRE is included in the PCRE distribution, and no further description is included here. The PCRE functions are called from Exim using the default option settings (that is, with no PCRE options set), except that the PCRE\_CASELESS option is set when the matching is required to be case-insensitive.

In most cases, when a regular expression is required in an Exim configuration, it has to start with a circumflex, in order to distinguish it from plain text or an “ends with” wildcard. In this example of a configuration setting, the second item in the colon-separated list is a regular expression.

```
domains = a.b.c : ^\\d{3} : *.y.z : ...
```

The doubling of the backslash is required because of string expansion that precedes interpretation – see section 11.1 for more discussion of this issue, and a way of avoiding the need for doubling backslashes. The regular expression that is eventually used in this example contains just one backslash. The circumflex is included in the regular expression, and has the normal effect of “anchoring” it to the start of the string that is being matched.

There are, however, two cases where a circumflex is not required for the recognition of a regular expression: these are the **match** condition in a string expansion, and the **matches** condition in an Exim filter file. In these cases, the relevant string is always treated as a regular expression; if it does not start with a circumflex, the expression is not anchored, and can match anywhere in the subject string.

In all cases, if you want a regular expression to match at the end of a string, you must code the \$ metacharacter to indicate this. For example:

```
domains = ^\\d{3}\\.example
```

matches the domain *123.example*, but it also matches *123.example.com*. You need to use:

```
domains = ^\\d{3}\\.example\\$
```

if you want *example* to be the top-level domain. The backslash before the \$ is needed because string expansion also interprets dollar characters.

## 9. File and database lookups

Exim can be configured to look up data in files or databases as it processes messages. Two different kinds of syntax are used:

- (1) A string that is to be expanded may contain explicit lookup requests. These cause parts of the string to be replaced by data that is obtained from the lookup. Lookups of this type are conditional expansion items. Different results can be defined for the cases of lookup success and failure. See chapter 11, where string expansions are described in detail. The key for the lookup is specified as part of the string expansion.
- (2) Lists of domains, hosts, and email addresses can contain lookup requests as a way of avoiding excessively long linear lists. In this case, the data that is returned by the lookup is often (but not always) discarded; whether the lookup succeeds or fails is what really counts. These kinds of list are described in chapter 10. The key for the lookup is given by the context in which the list is expanded.

String expansions, lists, and lookups interact with each other in such a way that there is no order in which to describe any one of them that does not involve references to the others. Each of these three chapters makes more sense if you have read the other two first. If you are reading this for the first time, be aware that some of it will make a lot more sense after you have read chapters 10 and 11.

### 9.1 Examples of different lookup syntax

It is easy to confuse the two different kinds of lookup, especially as the lists that may contain the second kind are always expanded before being processed as lists. Therefore, they may also contain lookups of the first kind. Be careful to distinguish between the following two examples:

```
domains = ${lookup{$sender_host_address}lsearch{/some/file}}
domains = lsearch;/some/file
```

The first uses a string expansion, the result of which must be a domain list. No strings have been specified for a successful or a failing lookup; the defaults in this case are the looked-up data and an empty string, respectively. The expansion takes place before the string is processed as a list, and the file that is searched could contain lines like this:

```
192.168.3.4: domain1:domain2:...
192.168.1.9: domain3:domain4:...
```

When the lookup succeeds, the result of the expansion is a list of domains (and possibly other types of item that are allowed in domain lists).

In the second example, the lookup is a single item in a domain list. It causes Exim to use a lookup to see if the domain that is being processed can be found in the file. The file could contain lines like this:

```
domain1:
domain2:
```

Any data that follows the key is not relevant when checking that the domain matches the list item.

It is possible, though no doubt confusing, to use both kinds of lookup at once. Consider a file containing lines like this:

```
192.168.5.6: lsearch;/another/file
```

If the value of `$sender_host_address` is 192.168.5.6, expansion of the first **domains** setting above generates the second setting, which therefore causes a second lookup to occur.

The rest of this chapter describes the different lookup types that are available. Any of them can be used in any part of the configuration where a lookup is permitted.

## 9.2 Lookup types

Two different types of data lookup are implemented:

- The *single-key* type requires the specification of a file in which to look, and a single key to search for. The key must be a non-empty string for the lookup to succeed. The lookup type determines how the file is searched.
- The *query-style* type accepts a generalized database query. No particular key value is assumed by Exim for query-style lookups. You can use whichever Exim variables you need to construct the database query.

The code for each lookup type is in a separate source file that is included in the binary of Exim only if the corresponding compile-time option is set. The default settings in *src/EDITME* are:

```
LOOKUP_DBM=yes
LOOKUP_LSEARCH=yes
```

which means that only linear searching and DBM lookups are included by default. For some types of lookup (e.g. SQL databases), you need to install appropriate libraries and header files before building Exim.

## 9.3 Single-key lookup types

The following single-key lookup types are implemented:

- *cdb*: The given file is searched as a Constant DataBase file, using the key string without a terminating binary zero. The cdb format is designed for indexed files that are read frequently and never updated, except by total re-creation. As such, it is particularly suitable for large files containing aliases or other indexed data referenced by an MTA. Information about cdb can be found in several places:

```
http://www.pobox.com/~djb/cdb.html
ftp://ftp.corpit.ru/pub/tinycdb/
http://packages.debian.org/stable/utils/freecdb.html
```

A cdb distribution is not needed in order to build Exim with cdb support, because the code for reading cdb files is included directly in Exim itself. However, no means of building or testing cdb files is provided with Exim, so you need to obtain a cdb distribution in order to do this.

- *dbm*: Calls to DBM library functions are used to extract data from the given DBM file by looking up the record with the given key. A terminating binary zero is included in the key that is passed to the DBM library. See section 4.4 for a discussion of DBM libraries.

For all versions of Berkeley DB, Exim uses the DB\_HASH style of database when building DBM files using the **exim\_dbmbuild** utility. However, when using Berkeley DB versions 3 or 4, it opens existing databases for reading with the DB\_UNKNOWN option. This enables it to handle any of the types of database that the library supports, and can be useful for accessing DBM files created by other applications. (For earlier DB versions, DB\_HASH is always used.)

- *dbmjz*: This is the same as *dbm*, except that the lookup key is interpreted as an Exim list; the elements of the list are joined together with ASCII NUL characters to form the lookup key. An example usage would be to authenticate incoming SMTP calls using the passwords from Cyrus SASL's */etc/sasl/db2* file with the *gsasl* authenticator or Exim's own *cram\_md5* authenticator.
- *dbmnz*: This is the same as *dbm*, except that a terminating binary zero is not included in the key that is passed to the DBM library. You may need this if you want to look up data in files that are created by or shared with some other application that does not use terminating zeros. For example, you need to use *dbmnz* rather than *dbm* if you want to authenticate incoming SMTP calls using the passwords from Courier's */etc/userdbshadow.dat* file. Exim's utility program for creating DBM files (*exim\_dbmbuild*) includes the zeros by default, but has an option to omit them (see section 53.9).
- *dsearch*: The given file must be a directory; this is searched for an entry whose name is the key by calling the *lstat()* function. The key may not contain any forward slash characters. If *lstat()* suc-

ceeds, the result of the lookup is the name of the entry, which may be a file, directory, symbolic link, or any other kind of directory entry. An example of how this lookup can be used to support virtual domains is given in section 50.7.

- *iplsearch*: The given file is a text file containing keys and data. A key is terminated by a colon or white space or the end of the line. The keys in the file must be IP addresses, or IP addresses with CIDR masks. Keys that involve IPv6 addresses must be enclosed in quotes to prevent the first internal colon being interpreted as a key terminator. For example:

```
1.2.3.4:          data for 1.2.3.4
192.168.0.0/16:   data for 192.168.0.0/16
"abcd::cdab":     data for abcd::cdab
"abcd:abcd::/32"  data for abcd:abcd::/32
```

The key for an *iplsearch* lookup must be an IP address (without a mask). The file is searched linearly, using the CIDR masks where present, until a matching key is found. The first key that matches is used; there is no attempt to find a “best” match. Apart from the way the keys are matched, the processing for *iplsearch* is the same as for *lsearch*.

**Warning 1:** Unlike most other single-key lookup types, a file of data for *iplsearch* can *not* be turned into a DBM or cdb file, because those lookup types support only literal keys.

**Warning 2:** In a host list, you must always use *net-iplsearch* so that the implicit key is the host’s IP address rather than its name (see section 10.12).

- *lsearch*: The given file is a text file that is searched linearly for a line beginning with the search key, terminated by a colon or white space or the end of the line. The search is case-insensitive; that is, upper and lower case letters are treated as the same. The first occurrence of the key that is found in the file is used.

White space between the key and the colon is permitted. The remainder of the line, with leading and trailing white space removed, is the data. This can be continued onto subsequent lines by starting them with any amount of white space, but only a single space character is included in the data at such a junction. If the data begins with a colon, the key must be terminated by a colon, for example:

```
baduser:  :fail:
```

Empty lines and lines beginning with # are ignored, even if they occur in the middle of an item. This is the traditional textual format of alias files. Note that the keys in an *lsearch* file are literal strings. There is no wildcarding of any kind.

In most *lsearch* files, keys are not required to contain colons or # characters, or white space. However, if you need this feature, it is available. If a key begins with a doublequote character, it is terminated only by a matching quote (or end of line), and the normal escaping rules apply to its contents (see section 6.17). An optional colon is permitted after quoted keys (exactly as for unquoted keys). There is no special handling of quotes for the data part of an *lsearch* line.

- *nis*: The given file is the name of a NIS map, and a NIS lookup is done with the given key, without a terminating binary zero. There is a variant called *nis0* which does include the terminating binary zero in the key. This is reportedly needed for Sun-style alias files. Exim does not recognize NIS aliases; the full map names must be used.
- *wildlsearch* or *nwildlsearch*: These search a file linearly, like *lsearch*, but instead of being interpreted as a literal string, each key in the file may be wildcarded. The difference between these two lookup types is that for *wildlsearch*, each key in the file is string-expanded before being used, whereas for *nwildlsearch*, no expansion takes place.

Like *lsearch*, the testing is done case-insensitively. However, keys in the file that are regular expressions can be made case-sensitive by the use of (-i) within the pattern. The following forms of wildcard are recognized:

- (1) The string may begin with an asterisk to mean “ends with”. For example:



```
*.a.b.c      data for anything.a.b.c
*fish       data for anythingfish
```

- (2) The string may begin with a circumflex to indicate a regular expression. For example, for *wildlsearch*:

```
^\N\d+\.a\.b\N    data for <digits>.a.b
```

Note the use of `\N` to disable expansion of the contents of the regular expression. If you are using *nwildlsearch*, where the keys are not string-expanded, the equivalent entry is:

```
^\d+\.a\.b        data for <digits>.a.b
```

The case-insensitive flag is set at the start of compiling the regular expression, but it can be turned off by using `(-i)` at an appropriate point. For example, to make the entire pattern case-sensitive:

```
^(?-i)\d+\.a\.b    data for <digits>.a.b
```

If the regular expression contains white space or colon characters, you must either quote it (see *lsearch* above), or represent these characters in other ways. For example, `\s` can be used for white space and `\x3A` for a colon. This may be easier than quoting, because if you quote, you have to escape all the backslashes inside the quotes.

**Note:** It is not possible to capture substrings in a regular expression match for later use, because the results of all lookups are cached. If a lookup is repeated, the result is taken from the cache, and no actual pattern matching takes place. The values of all the numeric variables are unset after a *(n)wildlsearch* match.

- (3) Although I cannot see it being of much use, the general matching function that is used to implement *(n)wildlsearch* means that the string may begin with a lookup name terminated by a semicolon, and followed by lookup data. For example:

```
cdb;/some/file    data for keys that match the file
```

The data that is obtained from the nested lookup is discarded.

Keys that do not match any of these patterns are interpreted literally. The continuation rules for the data are the same as for *lsearch*, and keys may be followed by optional colons.

**Warning:** Unlike most other single-key lookup types, a file of data for *(n)wildlsearch* can *not* be turned into a DBM or cdb file, because those lookup types support only literal keys.

## 9.4 Query-style lookup types

The supported query-style lookup types are listed below. Further details about many of them are given in later sections.

- *dnsdb*: This does a DNS search for one or more records whose domain names are given in the supplied query. The resulting data is the contents of the records. See section 9.10.
- *ibase*: This does a lookup in an InterBase database.
- *ldap*: This does an LDAP lookup using a query in the form of a URL, and returns attributes from a single entry. There is a variant called *ldapm* that permits values from multiple entries to be returned. A third variant called *ldapdn* returns the Distinguished Name of a single entry instead of any attribute values. See section 9.14.
- *mysql*: The format of the query is an SQL statement that is passed to a MySQL database. See section 9.21.
- *nisplus*: This does a NIS+ lookup using a query that can specify the name of the field to be returned. See section 9.20.
- *oracle*: The format of the query is an SQL statement that is passed to an Oracle database. See section 9.21.

- *passwd* is a query-style lookup with queries that are just user names. The lookup calls *getpwnam()* to interrogate the system password data, and on success, the result string is the same as you would get from an *lsearch* lookup on a traditional */etc/passwd* file, though with *\** for the password value. For example:

```
*:42:42:King Rat:/home/kr:/bin/bash
```

- *pgsql*: The format of the query is an SQL statement that is passed to a PostgreSQL database. See section 9.21.
- *redis*: The format of the query is either a simple get or simple set, passed to a Redis database. See section 9.21.
- *sqlite*: The format of the query is a file name followed by an SQL statement that is passed to an SQLite database. See section 9.26.
- *testdb*: This is a lookup type that is used for testing Exim. It is not likely to be useful in normal operation.
- *whoson*: *Whoson* (<http://whoson.sourceforge.net>) is a protocol that allows a server to check whether a particular (dynamically allocated) IP address is currently allocated to a known (trusted) user and, optionally, to obtain the identity of the said user. For SMTP servers, *Whoson* was popular at one time for “POP before SMTP” authentication, but that approach has been superseded by SMTP authentication. In Exim, *Whoson* can be used to implement “POP before SMTP” checking using ACL statements such as

```
require condition = \
    ${lookup whoson {$sender_host_address}{yes}{no}}
```

The query consists of a single IP address. The value returned is the name of the authenticated user, which is stored in the variable *\$value*. However, in this example, the data in *\$value* is not used; the result of the lookup is one of the fixed strings “yes” or “no”.

## 9.5 Temporary errors in lookups

Lookup functions can return temporary error codes if the lookup cannot be completed. For example, an SQL or LDAP database might be unavailable. For this reason, it is not advisable to use a lookup that might do this for critical options such as a list of local domains.

When a lookup cannot be completed in a router or transport, delivery of the message (to the relevant address) is deferred, as for any other temporary error. In other circumstances Exim may assume the lookup has failed, or may give up altogether.

## 9.6 Default values in single-key lookups

In this context, a “default value” is a value specified by the administrator that is to be used if a lookup fails.

**Note:** This section applies only to single-key lookups. For query-style lookups, the facilities of the query language must be used. An attempt to specify a default for a query-style lookup provokes an error.

If “\*” is added to a single-key lookup type (for example, **lsearch\***) and the initial lookup fails, the key “\*” is looked up in the file to provide a default value. See also the section on partial matching below.

Alternatively, if “\*@” is added to a single-key lookup type (for example **dbm\*@**) then, if the initial lookup fails and the key contains an @ character, a second lookup is done with everything before the last @ replaced by \*. This makes it possible to provide per-domain defaults in alias files that include the domains in the keys. If the second lookup fails (or doesn’t take place because there is no @ in the key), “\*” is looked up. For example, a *redirect* router might contain:

```
data = ${lookup{$local_part@$domain}lsearch*{/etc/mix-aliases}}
```

Suppose the address that is being processed is *jane@eyre.example*. Exim looks up these keys, in this order:

```
jane@eyre.example
*@eyre.example
*
```

The data is taken from whichever key it finds first. **Note:** In an *lsearch* file, this does not mean the first of these keys in the file. A complete scan is done for each key, and only if it is not found at all does Exim move on to try the next key.

## 9.7 Partial matching in single-key lookups

The normal operation of a single-key lookup is to search the file for an exact match with the given key. However, in a number of situations where domains are being looked up, it is useful to be able to do partial matching. In this case, information in the file that has a key starting with “\*.” is matched by any domain that ends with the components that follow the full stop. For example, if a key in a DBM file is

```
*.dates.fict.example
```

then when partial matching is enabled this is matched by (amongst others) *2001.dates.fict.example* and *1984.dates.fict.example*. It is also matched by *dates.fict.example*, if that does not appear as a separate key in the file.

**Note:** Partial matching is not available for query-style lookups. It is also not available for any lookup items in address lists (see section 10.19).

Partial matching is implemented by doing a series of separate lookups using keys constructed by modifying the original subject key. This means that it can be used with any of the single-key lookup types, provided that partial matching keys beginning with a special prefix (default “\*.”) are included in the data file. Keys in the file that do not begin with the prefix are matched only by unmodified subject keys when partial matching is in use.

Partial matching is requested by adding the string “partial-” to the front of the name of a single-key lookup type, for example, **partial-dbm**. When this is done, the subject key is first looked up unmodified; if that fails, “\*.” is added at the start of the subject key, and it is looked up again. If that fails, further lookups are tried with dot-separated components removed from the start of the subject key, one-by-one, and “\*.” added on the front of what remains.

A minimum number of two non-\* components are required. This can be adjusted by including a number before the hyphen in the search type. For example, **partial3-lsearch** specifies a minimum of three non-\* components in the modified keys. Omitting the number is equivalent to “partial2-”. If the subject key is *2250.dates.fict.example* then the following keys are looked up when the minimum number of non-\* components is two:

```
2250.dates.fict.example
*.2250.dates.fict.example
*.dates.fict.example
*.fict.example
```

As soon as one key in the sequence is successfully looked up, the lookup finishes.

The use of “\*.” as the partial matching prefix is a default that can be changed. The motivation for this feature is to allow Exim to operate with file formats that are used by other MTAs. A different prefix can be supplied in parentheses instead of the hyphen after “partial”. For example:

```
domains = partial(.)lsearch;/some/file
```

In this example, if the domain is *a.b.c*, the sequence of lookups is *a.b.c*, *.a.b.c*, and *.b.c* (the default minimum of 2 non-wild components is unchanged). The prefix may consist of any punctuation characters other than a closing parenthesis. It may be empty, for example:

```
domains = partial1()cdb;/some/file
```

For this example, if the domain is *a.b.c*, the sequence of lookups is *a.b.c*, *b.c*, and *c*.

If “partial0” is specified, what happens at the end (when the lookup with just one non-wild component has failed, and the original key is shortened right down to the null string) depends on the prefix:

- If the prefix has zero length, the whole lookup fails.
- If the prefix has length 1, a lookup for just the prefix is done. For example, the final lookup for “partial0(.”) is for . alone.
- Otherwise, if the prefix ends in a dot, the dot is removed, and the remainder is looked up. With the default prefix, therefore, the final lookup is for “\*” on its own.
- Otherwise, the whole prefix is looked up.

If the search type ends in “\*” or “\*@” (see section 9.6 above), the search for an ultimate default that this implies happens after all partial lookups have failed. If “partial0” is specified, adding “\*” to the search type has no effect with the default prefix, because the “\*” key is already included in the sequence of partial lookups. However, there might be a use for lookup types such as “partial0(.)lsearch\*”.

The use of “\*” in lookup partial matching differs from its use as a wildcard in domain lists and the like. Partial matching works only in terms of dot-separated components; a key such as \*fict.example in a database file is useless, because the asterisk in a partial matching subject key is always followed by a dot.

## 9.8 Lookup caching

Exim caches all lookup results in order to avoid needless repetition of lookups. However, because (apart from the daemon) Exim operates as a collection of independent, short-lived processes, this caching applies only within a single Exim process. There is no inter-process lookup caching facility.

For single-key lookups, Exim keeps the relevant files open in case there is another lookup that needs them. In some types of configuration this can lead to many files being kept open for messages with many recipients. To avoid hitting the operating system limit on the number of simultaneously open files, Exim closes the least recently used file when it needs to open more files than its own internal limit, which can be changed via the **lookup\_open\_max** option.

The single-key lookup files are closed and the lookup caches are flushed at strategic points during delivery – for example, after all routing is complete.

## 9.9 Quoting lookup data

When data from an incoming message is included in a query-style lookup, there is the possibility of special characters in the data messing up the syntax of the query. For example, a NIS+ query that contains

```
[name=$local_part]
```

will be broken if the local part happens to contain a closing square bracket. For NIS+, data can be enclosed in double quotes like this:

```
[name="$local_part"]
```

but this still leaves the problem of a double quote in the data. The rule for NIS+ is that double quotes must be doubled. Other lookup types have different rules, and to cope with the differing requirements, an expansion operator of the following form is provided:

```
${quote_<lookup-type>:<string>}
```

For example, the safest way to write the NIS+ query is

```
[name="${quote_nisplus:$local_part}"]
```

See chapter 11 for full coverage of string expansions. The quote operator can be used for all lookup types, but has no effect for single-key lookups, since no quoting is ever needed in their key strings.

## 9.10 More about dnsdb

The *dnsdb* lookup type uses the DNS as its database. A simple query consists of a record type and a domain name, separated by an equals sign. For example, an expansion string could contain:

```
${lookup dnsdb{mx=a.b.example}}{$value}fail}
```

If the lookup succeeds, the result is placed in *\$value*, which in this case is used on its own as the result. If the lookup does not succeed, the *fail* keyword causes a *forced expansion failure* – see section 11.4 for an explanation of what this means.

The supported DNS record types are A, CNAME, MX, NS, PTR, SOA, SPF, SRV, TLSA and TXT, and, when Exim is compiled with IPv6 support, AAAA. If no type is given, TXT is assumed.

For any record type, if multiple records are found, the data is returned as a concatenation, with newline as the default separator. The order, of course, depends on the DNS resolver. You can specify a different separator character between multiple records by putting a right angle-bracket followed immediately by the new separator at the start of the query. For example:

```
${lookup dnsdb{>: a=host1.example}}
```

It is permitted to specify a space as the separator character. Further white space is ignored. For lookup types that return multiple fields per record, an alternate field separator can be specified using a comma after the main separator character, followed immediately by the field separator.

When the type is PTR, the data can be an IP address, written as normal; inversion and the addition of **in-addr.arpa** or **ip6.arpa** happens automatically. For example:

```
${lookup dnsdb{ptr=192.168.4.5}}{$value}fail}
```

If the data for a PTR record is not a syntactically valid IP address, it is not altered and nothing is added.

For an MX lookup, both the preference value and the host name are returned for each record, separated by a space. For an SRV lookup, the priority, weight, port, and host name are returned for each record, separated by spaces. The field separator can be modified as above.

For TXT records with multiple items of data, only the first item is returned, unless a field separator is specified. To concatenate items without a separator, use a semicolon instead. For SPF records the default behaviour is to concatenate multiple items without using a separator.

```
${lookup dnsdb{>\n, : txt=a.b.example}}
${lookup dnsdb{>\n; txt=a.b.example}}
${lookup dnsdb{spf=example.org}}
```

It is permitted to specify a space as the separator character. Further white space is ignored.

For an SOA lookup, while no result is obtained the lookup is redone with successively more leading components dropped from the given domain. Only the primary-nameserver field is returned unless a field separator is specified.

```
${lookup dnsdb{>:, ; soa=a.b.example.com}}
```

## 9.11 Dnsdb lookup modifiers

Modifiers for *dnsdb* lookups are given by optional keywords, each followed by a comma, that may appear before the record type.

The *dnsdb* lookup fails only if all the DNS lookups fail. If there is a temporary DNS error for any of them, the behaviour is controlled by a defer-option modifier. The possible keywords are “defer\_strict”, “defer\_never”, and “defer\_lax”. With “strict” behaviour, any temporary DNS error causes the whole lookup to defer. With “never” behaviour, a temporary DNS error is ignored, and the behaviour is as if the DNS lookup failed to find anything. With “lax” behaviour, all the queries are attempted, but a temporary DNS error causes the whole lookup to defer only if none of the other lookups succeed. The default is “lax”, so the following lookups are equivalent:

```

${lookup dnsdb{defer_lax,a=one.host.com:two.host.com}}
${lookup dnsdb{a=one.host.com:two.host.com}}

```

Thus, in the default case, as long as at least one of the DNS lookups yields some data, the lookup succeeds.

Use of *DNSSEC* is controlled by a *dnssec* modifier. The possible keywords are “*dnssec\_strict*”, “*dnssec\_lax*”, and “*dnssec\_never*”. With “*strict*” or “*lax*” DNSSEC information is requested with the lookup. With “*strict*” a response from the DNS resolver that is not labelled as authenticated data is treated as equivalent to a temporary DNS error. The default is “*never*”.

See also the *\$lookup\_dnssec\_authenticated* variable.

Timeout for the *dnsdb* lookup can be controlled by a *retrans* modifier. The form is “*retrans\_VAL*” where *VAL* is an Exim time specification (e.g. “5s”). The default value is set by the main configuration option ***dns\_retrans***.

Retries for the *dnsdb* lookup can be controlled by a *retry* modifier. The form is “*retry\_VAL*” where *VAL* is an integer. The default count is set by the main configuration option ***dns\_retry***.

*Dnsdb* lookup results are cached within a single process (and its children). The cache entry lifetime is limited to the smallest time-to-live (TTL) value of the set of returned DNS records.

## 9.12 Pseudo *dnsdb* record types

By default, both the preference value and the host name are returned for each MX record, separated by a space. If you want only host names, you can use the pseudo-type MXH:

```

${lookup dnsdb{mxh=a.b.example}}

```

In this case, the preference values are omitted, and just the host names are returned.

Another pseudo-type is ZNS (for “zone NS”). It performs a lookup for NS records on the given domain, but if none are found, it removes the first component of the domain name, and tries again. This process continues until NS records are found or there are no more components left (or there is a DNS error). In other words, it may return the name servers for a top-level domain, but it never returns the root name servers. If there are no NS records for the top-level domain, the lookup fails. Consider these examples:

```

${lookup dnsdb{zns=xxx.quercite.com}}
${lookup dnsdb{zns=xxx.edu}}

```

Assuming that in each case there are no NS records for the full domain name, the first returns the name servers for **quercite.com**, and the second returns the name servers for **edu**.

You should be careful about how you use this lookup because, unless the top-level domain does not exist, the lookup always returns some host names. The sort of use to which this might be put is for seeing if the name servers for a given domain are on a blacklist. You can probably assume that the name servers for the high-level domains such as **com** or **co.uk** are not going to be on such a list.

A third pseudo-type is CSA (Client SMTP Authorization). This looks up SRV records according to the CSA rules, which are described in section 43.50. Although *dnsdb* supports SRV lookups directly, this is not sufficient because of the extra parent domain search behaviour of CSA. The result of a successful lookup such as:

```

${lookup dnsdb {csa=$sender_helo_name}}

```

has two space-separated fields: an authorization code and a target host name. The authorization code can be “Y” for yes, “N” for no, “X” for explicit authorization required but absent, or “?” for unknown.

The pseudo-type A+ performs an AAAA and then an A lookup. All results are returned; defer processing (see below) is handled separately for each lookup. Example:

```

${lookup dnsdb {>; a+=$sender_helo_name}}

```

## 9.13 Multiple dnsdb lookups

In the previous sections, *dnsdb* lookups for a single domain are described. However, you can specify a list of domains or IP addresses in a single *dnsdb* lookup. The list is specified in the normal Exim way, with colon as the default separator, but with the ability to change this. For example:

```
{lookup dnsdb{one.domain.com:two.domain.com}}
{lookup dnsdb{a=one.host.com:two.host.com}}
{lookup dnsdb{ptr = <; 1.2.3.4 ; 4.5.6.8}}
```

In order to retain backwards compatibility, there is one special case: if the lookup type is PTR and no change of separator is specified, Exim looks to see if the rest of the string is precisely one IPv6 address. In this case, it does not treat it as a list.

The data from each lookup is concatenated, with newline separators by default, in the same way that multiple DNS records for a single item are handled. A different separator can be specified, as described above.

## 9.14 More about LDAP

The original LDAP implementation came from the University of Michigan; this has become “Open LDAP”, and there are now two different releases. Another implementation comes from Netscape, and Solaris 7 and subsequent releases contain inbuilt LDAP support. Unfortunately, though these are all compatible at the lookup function level, their error handling is different. For this reason it is necessary to set a compile-time variable when building Exim with LDAP, to indicate which LDAP library is in use. One of the following should appear in your *Local/Makefile*:

```
LDAP_LIB_TYPE=UMICHIGAN
LDAP_LIB_TYPE=OPENLDAP1
LDAP_LIB_TYPE=OPENLDAP2
LDAP_LIB_TYPE=NETSCAPE
LDAP_LIB_TYPE=SOLARIS
```

If `LDAP_LIB_TYPE` is not set, Exim assumes `OPENLDAP1`, which has the same interface as the University of Michigan version.

There are three LDAP lookup types in Exim. These behave slightly differently in the way they handle the results of a query:

- *ldap* requires the result to contain just one entry; if there are more, it gives an error.
- *ldapdn* also requires the result to contain just one entry, but it is the Distinguished Name that is returned rather than any attribute values.
- *ldapm* permits the result to contain more than one entry; the attributes from all of them are returned.

For *ldap* and *ldapm*, if a query finds only entries with no attributes, Exim behaves as if the entry did not exist, and the lookup fails. The format of the data returned by a successful lookup is described in the next section. First we explain how LDAP queries are coded.

## 9.15 Format of LDAP queries

An LDAP query takes the form of a URL as defined in RFC 2255. For example, in the configuration of a *redirect* router one might have this setting:

```
data = ${lookup ldap \
  {ldap:///cn=$local_part,o=University%20of%20Cambridge,\
  c=UK?mailbox?base?}}
```

The URL may begin with `ldap` or `ldaps` if your LDAP library supports secure (encrypted) LDAP connections. The second of these ensures that an encrypted TLS connection is used.

With sufficiently modern LDAP libraries, Exim supports forcing TLS over regular LDAP connections, rather than the SSL-on-connect `ldaps`. See the **`ldap_start_tls`** option.

Starting with Exim 4.83, the initialization of LDAP with TLS is more tightly controlled. Every part of the TLS configuration can be configured by settings in *exim.conf*. Depending on the version of the client libraries installed on your system, some of the initialization may have required setting options in */etc/ldap.conf* or *~/ldaprc* to get TLS working with self-signed certificates. This revealed a nuance where the current UID that exim was running as could affect which config files it read. With Exim 4.83, these methods become optional, only taking effect if not specifically set in *exim.conf*.

## 9.16 LDAP quoting

Two levels of quoting are required in LDAP queries, the first for LDAP itself and the second because the LDAP query is represented as a URL. Furthermore, within an LDAP query, two different kinds of quoting are required. For this reason, there are two different LDAP-specific quoting operators.

The **quote\_ldap** operator is designed for use on strings that are part of filter specifications. Conceptually, it first does the following conversions on the string:

```
*    =>    \2A
(    =>    \28
)    =>    \29
\    =>    \5C
```

in accordance with RFC 2254. The resulting string is then quoted according to the rules for URLs, that is, all non-alphanumeric characters except

```
! $ ' - . _ ( ) * +
```

are converted to their hex values, preceded by a percent sign. For example:

```
${quote_ldap: a(bc)*, a<yz>; }
```

yields

```
%20a%5C28bc%5C29%5C2A%2C%20a%3Cyz%3E%3B%20
```

Removing the URL quoting, this is (with a leading and a trailing space):

```
a\28bc\29\2A, a<yz>;
```

The **quote\_ldap\_dn** operator is designed for use on strings that are part of base DN specifications in queries. Conceptually, it first converts the string by inserting a backslash in front of any of the following characters:

```
, + " \ < > ;
```

It also inserts a backslash before any leading spaces or # characters, and before any trailing spaces. (These rules are in RFC 2253.) The resulting string is then quoted according to the rules for URLs. For example:

```
${quote_ldap_dn: a(bc)*, a<yz>; }
```

yields

```
%5C%20a(bc)*%5C%2C%20a%5C%3Cyz%5C%3E%5C%3B%5C%20
```

Removing the URL quoting, this is (with a trailing space):

```
\ a(bc)*\, a<yz>\>\;
```

There are some further comments about quoting in the section on LDAP authentication below.

## 9.17 LDAP connections

The connection to an LDAP server may either be over TCP/IP, or, when OpenLDAP is in use, via a Unix domain socket. The example given above does not specify an LDAP server. A server that is reached by TCP/IP can be specified in a query by starting it with

```
ldap://<hostname>:<port>/...
```



If the port (and preceding colon) are omitted, the standard LDAP port (389) is used. When no server is specified in a query, a list of default servers is taken from the **ldap\_default\_servers** configuration option. This supplies a colon-separated list of servers which are tried in turn until one successfully handles a query, or there is a serious error. Successful handling either returns the requested data, or indicates that it does not exist. Serious errors are syntactical, or multiple values when only a single value is expected. Errors which cause the next server to be tried are connection failures, bind failures, and timeouts.

For each server name in the list, a port number can be given. The standard way of specifying a host and port is to use a colon separator (RFC 1738). Because **ldap\_default\_servers** is a colon-separated list, such colons have to be doubled. For example

```
ldap_default_servers = ldap1.example.com::145:ldap2.example.com
```

If **ldap\_default\_servers** is unset, a URL with no server name is passed to the LDAP library with no server name, and the library's default (normally the local host) is used.

If you are using the OpenLDAP library, you can connect to an LDAP server using a Unix domain socket instead of a TCP/IP connection. This is specified by using `ldapi` instead of `ldap` in LDAP queries. What follows here applies only to OpenLDAP. If Exim is compiled with a different LDAP library, this feature is not available.

For this type of connection, instead of a host name for the server, a pathname for the socket is required, and the port number is not relevant. The pathname can be specified either as an item in **ldap\_default\_servers**, or inline in the query. In the former case, you can have settings such as

```
ldap_default_servers = /tmp/ldap.sock : backup.ldap.your.domain
```

When the pathname is given in the query, you have to escape the slashes as `%2F` to fit in with the LDAP URL syntax. For example:

```
${lookup ldap {ldapi://%2Ftmp%2Fldap.sock/o=...
```

When Exim processes an LDAP lookup and finds that the “hostname” is really a pathname, it uses the Unix domain socket code, even if the query actually specifies `ldap` or `ldaps`. In particular, no encryption is used for a socket connection. This behaviour means that you can use a setting of **ldap\_default\_servers** such as in the example above with traditional `ldap` or `ldaps` queries, and it will work. First, Exim tries a connection via the Unix domain socket; if that fails, it tries a TCP/IP connection to the backup host.

If an explicit `ldapi` type is given in a query when a host name is specified, an error is diagnosed. However, if there are more items in **ldap\_default\_servers**, they are tried. In other words:

- Using a pathname with `ldap` or `ldaps` forces the use of the Unix domain interface.
- Using `ldapi` with a host name causes an error.

Using `ldapi` with no host or path in the query, and no setting of **ldap\_default\_servers**, does whatever the library does by default.

## 9.18 LDAP authentication and control information

The LDAP URL syntax provides no way of passing authentication and other control information to the server. To make this possible, the URL in an LDAP query may be preceded by any number of `<name>=<value>` settings, separated by spaces. If a value contains spaces it must be enclosed in double quotes, and when double quotes are used, backslash is interpreted in the usual way inside them. The following names are recognized:

DEREFERENCE	set the dereferencing parameter
NETTIME	set a timeout for a network operation
USER	set the DN, for authenticating the LDAP bind
PASS	set the password, likewise
REFERRALS	set the referrals parameter
SERVERS	set alternate server list for this query only

SIZE	set the limit for the number of entries returned
TIME	set the maximum waiting time for a query

The value of the DEREERENCE parameter must be one of the words “never”, “searching”, “finding”, or “always”. The value of the REFERRALS parameter must be “follow” (the default) or “nofollow”. The latter stops the LDAP library from trying to follow referrals issued by the LDAP server.

The name CONNECT is an obsolete name for NETTIME, retained for backwards compatibility. This timeout (specified as a number of seconds) is enforced from the client end for operations that can be carried out over a network. Specifically, it applies to network connections and calls to the *ldap\_result()* function. If the value is greater than zero, it is used if LDAP\_OPT\_NETWORK\_TIMEOUT is defined in the LDAP headers (OpenLDAP), or if LDAP\_X\_OPT\_CONNECT\_TIMEOUT is defined in the LDAP headers (Netscape SDK 4.1). A value of zero forces an explicit setting of “no timeout” for Netscape SDK; for OpenLDAP no action is taken.

The TIME parameter (also a number of seconds) is passed to the server to set a server-side limit on the time taken to complete a search.

The SERVERS parameter allows you to specify an alternate list of ldap servers to use for an individual lookup. The global **ldap\_default\_servers** option provides a default list of ldap servers, and a single lookup can specify a single ldap server to use. But when you need to do a lookup with a list of servers that is different than the default list (maybe different order, maybe a completely different set of servers), the SERVERS parameter allows you to specify this alternate list (colon-separated).

Here is an example of an LDAP query in an Exim lookup that uses some of these values. This is a single line, folded to fit on the page:

```
{lookup ldap
  {user="cn=manager,o=University of Cambridge,c=UK" pass=secret
  ldap:///o=University%20of%20Cambridge,c=UK?sn?sub?(cn=foo) }
  {$value}fail}
```

The encoding of spaces as %20 is a URL thing which should not be done for any of the auxiliary data. Exim configuration settings that include lookups which contain password information should be preceded by “hide” to prevent non-admin users from using the **-bP** option to see their values.

The auxiliary data items may be given in any order. The default is no connection timeout (the system timeout is used), no user or password, no limit on the number of entries returned, and no time limit on queries.

When a DN is quoted in the USER= setting for LDAP authentication, Exim removes any URL quoting that it may contain before passing it LDAP. Apparently some libraries do this for themselves, but some do not. Removing the URL quoting has two advantages:

- It makes it possible to use the same **quote\_ldap\_dn** expansion for USER= DN's as with DN's inside actual queries.
- It permits spaces inside USER= DN's.

For example, a setting such as

```
USER=cn=${quote_ldap_dn:$1}
```

should work even if *\$1* contains spaces.

Expanded data for the PASS= value should be quoted using the **quote** expansion operator, rather than the LDAP quote operators. The only reason this field needs quoting is to ensure that it conforms to the Exim syntax, which does not allow unquoted spaces. For example:

```
PASS=${quote:$3}
```

The LDAP authentication mechanism can be used to check passwords as part of SMTP authentication. See the **ldapauth** expansion string condition in chapter 11.

## 9.19 Format of data returned by LDAP

The *ldapdn* lookup type returns the Distinguished Name from a single entry as a sequence of values, for example

```
cn=manager,o=University of Cambridge,c=UK
```

The *ldap* lookup type generates an error if more than one entry matches the search filter, whereas *ldapm* permits this case, and inserts a newline in the result between the data from different entries. It is possible for multiple values to be returned for both *ldap* and *ldapm*, but in the former case you know that whatever values are returned all came from a single entry in the directory.

In the common case where you specify a single attribute in your LDAP query, the result is not quoted, and does not contain the attribute name. If the attribute has multiple values, they are separated by commas. Any comma that is part of an attribute's value is doubled.

If you specify multiple attributes, the result contains space-separated, quoted strings, each preceded by the attribute name and an equals sign. Within the quotes, the quote character, backslash, and newline are escaped with backslashes, and commas are used to separate multiple values for the attribute. Any commas in attribute values are doubled (permitting treatment of the values as a comma-separated list). Apart from the escaping, the string within quotes takes the same form as the output when a single attribute is requested. Specifying no attributes is the same as specifying all of an entry's attributes.

Here are some examples of the output format. The first line of each pair is an LDAP query, and the second is the data that is returned. The attribute called **attr1** has two values, one of them with an embedded comma, whereas **attr2** has only one value. Both attributes are derived from **attr** (they have SUP **attr** in their schema definitions).

```
ldap:///o=base?attr1?sub?(uid=fred)
value1.1,value1,,2
```

```
ldap:///o=base?attr2?sub?(uid=fred)
value two
```

```
ldap:///o=base?attr?sub?(uid=fred)
value1.1,value1,,2,value two
```

```
ldap:///o=base?attr1,attr2?sub?(uid=fred)
attr1="value1.1,value1,,2" attr2="value two"
```

```
ldap:///o=base??sub?(uid=fred)
objectClass="top" attr1="value1.1,value1,,2" attr2="value two"
```

You can make use of Exim's **-be** option to run expansion tests and thereby check the results of LDAP lookups. The **extract** operator in string expansions can be used to pick out individual fields from data that consists of *key=value* pairs. The **listextract** operator should be used to pick out individual values of attributes, even when only a single value is expected. The doubling of embedded commas allows you to use the returned data as a comma separated list (using the "<," syntax for changing the input list separator).

## 9.20 More about NIS+

NIS+ queries consist of a NIS+ *indexed name* followed by an optional colon and field name. If this is given, the result of a successful query is the contents of the named field; otherwise the result consists of a concatenation of *field-name=field-value* pairs, separated by spaces. Empty values and values containing spaces are quoted. For example, the query

```
[name=mg1456],passwd.org_dir
```

might return the string

```
name=mg1456 passwd="" uid=999 gid=999 gcos="Martin Guerre"
home=/home/mg1456 shell=/bin/bash shadow=""
```

(split over two lines here to fit on the page), whereas

```
[name=mg1456],passwd.org_dir:gcos
```

would just return

```
Martin Guerre
```

with no quotes. A NIS+ lookup fails if NIS+ returns more than one table entry for the given indexed key. The effect of the **quote\_nisplus** expansion operator is to double any quote characters within the text.

## 9.21 SQL lookups

Exim can support lookups in InterBase, MySQL, Oracle, PostgreSQL, Redis, and SQLite databases. Queries for these databases contain SQL statements, so an example might be

```
${lookup mysql{select mailbox from users where id='userx'}\
{$value}fail}
```

If the result of the query contains more than one field, the data for each field in the row is returned, preceded by its name, so the result of

```
${lookup pgsql{select home,name from users where id='userx'}\
{$value}}
```

might be

```
home=/home/userx name="Mister X"
```

Empty values and values containing spaces are double quoted, with embedded quotes escaped by a backslash. If the result of the query contains just one field, the value is passed back verbatim, without a field name, for example:

```
Mister X
```

If the result of the query yields more than one row, it is all concatenated, with a newline between the data for each row.

## 9.22 More about MySQL, PostgreSQL, Oracle, InterBase, and Redis

If any MySQL, PostgreSQL, Oracle, InterBase or Redis lookups are used, the **mysql\_servers**, **pgsql\_servers**, **oracle\_servers**, **ibase\_servers**, or **redis\_servers** option (as appropriate) must be set to a colon-separated list of server information. (For MySQL and PostgreSQL, the global option need not be set if all queries contain their own server information – see section 9.23.) For all but Redis each item in the list is a slash-separated list of four items: host name, database name, user name, and password. In the case of Oracle, the host name field is used for the “service name”, and the database name field is not used and should be empty. For example:

```
hide oracle_servers = oracle.plc.example//userx/abcdwxyz
```

Because password data is sensitive, you should always precede the setting with “hide”, to prevent non-admin users from obtaining the setting via the **-bP** option. Here is an example where two MySQL servers are listed:

```
hide mysql_servers = localhost/users/root/secret:\
otherhost/users/root/othersecret
```

For MySQL and PostgreSQL, a host may be specified as **<name>:<port>** but because this is a colon-separated list, the colon has to be doubled. For each query, these parameter groups are tried in order until a connection is made and a query is successfully processed. The result of a query may be that no data is found, but that is still a successful query. In other words, the list of servers provides a backup facility, not a list of different places to look.

For Redis the global option need not be specified if all queries contain their own server information – see section 9.23. If specified, the option must be set to a colon-separated list of server information. Each item in the list is a slash-separated list of three items: host, database number, and password.

- (1) The host is required and may be either an IPv4 address and optional port number (separated by a colon, which needs doubling due to the higher-level list), or a Unix socket pathname enclosed in parentheses
- (2) The database number is optional; if present that number is selected in the backend
- (3) The password is optional; if present it is used to authenticate to the backend

The **quote\_mysql**, **quote\_pgsql**, and **quote\_oracle** expansion operators convert newline, tab, carriage return, and backspace to `\n`, `\t`, `\r`, and `\b` respectively, and the characters single-quote, double-quote, and backslash itself are escaped with backslashes.

The **quote\_redis** expansion operator escapes whitespace and backslash characters with a backslash.

## 9.23 Specifying the server in the query

For MySQL, PostgreSQL and Redis lookups (but not currently for Oracle and InterBase), it is possible to specify a list of servers with an individual query. This is done by starting the query with

```
servers=server1:server2:server3:...
```

Each item in the list may take one of two forms:

- (1) If it contains no slashes it is assumed to be just a host name. The appropriate global option (**mysql\_servers** or **pgsql\_servers**) is searched for a host of the same name, and the remaining parameters (database, user, password) are taken from there.
- (2) If it contains any slashes, it is taken as a complete parameter set.

The list of servers is used in exactly the same way as the global list. Once a connection to a server has happened and a query has been successfully executed, processing of the lookup ceases.

This feature is intended for use in master/slave situations where updates are occurring and you want to update the master rather than a slave. If the master is in the list as a backup for reading, you might have a global setting like this:

```
mysql_servers = slave1/db/name/pw:\
               slave2/db/name/pw:\
               master/db/name/pw
```

In an updating lookup, you could then write:

```
${lookup mysql{servers=master; UPDATE ...} }
```

That query would then be sent only to the master server. If, on the other hand, the master is not to be used for reading, and so is not present in the global option, you can still update it by a query of this form:

```
${lookup pgsql{servers=master/db/name/pw; UPDATE ...} }
```

## 9.24 Special MySQL features

For MySQL, an empty host name or the use of “localhost” in **mysql\_servers** causes a connection to the server on the local host by means of a Unix domain socket. An alternate socket can be specified in parentheses. An option group name for MySQL option files can be specified in square brackets; the default value is “exim”. The full syntax of each item in **mysql\_servers** is:

```
<hostname>::<port>(<socket name>)[<option group>]/<database>/<user>/<password>
```

Any of the four sub-parts of the first field can be omitted. For normal use on the local host it can be left blank or set to just “localhost”.

No database need be supplied – but if it is absent here, it must be given in the queries.

If a MySQL query is issued that does not request any data (an insert, update, or delete command), the result of the lookup is the number of rows affected.

**Warning:** This can be misleading. If an update does not actually change anything (for example, setting a field to the value it already has), the result is zero because no rows are affected.

## 9.25 Special PostgreSQL features

PostgreSQL lookups can also use Unix domain socket connections to the database. This is usually faster and costs less CPU time than a TCP/IP connection. However it can be used only if the mail server runs on the same machine as the database server. A configuration line for PostgreSQL via Unix domain sockets looks like this:

```
hide pgsql_servers = (/tmp/.s.PGSQL.5432)/db/user/password : ...
```

In other words, instead of supplying a host name, a path to the socket is given. The path name is enclosed in parentheses so that its slashes aren't visually confused with the delimiters for the other server parameters.

If a PostgreSQL query is issued that does not request any data (an insert, update, or delete command), the result of the lookup is the number of rows affected.

## 9.26 More about SQLite

SQLite is different to the other SQL lookups because a file name is required in addition to the SQL query. An SQLite database is a single file, and there is no daemon as in the other SQL databases. The interface to Exim requires the name of the file, as an absolute path, to be given at the start of the query. It is separated from the query by white space. This means that the path name cannot contain white space. Here is a lookup expansion example:

```
${lookup sqlite {/some/thing/sqlitedb \
  select name from aliases where id='userx';}}
```

In a list, the syntax is similar. For example:

```
domainlist relay_to_domains = sqlite;/some/thing/sqlitedb \
  select * from relays where ip='$sender_host_address';
```

The only character affected by the **quote\_sqlite** operator is a single quote, which it doubles.

The SQLite library handles multiple simultaneous accesses to the database internally. Multiple readers are permitted, but only one process can update at once. Attempts to access the database while it is being updated are rejected after a timeout period, during which the SQLite library waits for the lock to be released. In Exim, the default timeout is set to 5 seconds, but it can be changed by means of the **sqlite\_lock\_timeout** option.

## 9.27 More about Redis

Redis is a non-SQL database. Commands are simple get and set. Examples:

```
${lookup redis{set keyname ${quote_redis:objvalue plus}}}  
${lookup redis{get keyname}}
```

## 10. Domain, host, address, and local part lists

A number of Exim configuration options contain lists of domains, hosts, email addresses, or local parts. For example, the **hold\_domains** option contains a list of domains whose delivery is currently suspended. These lists are also used as data in ACL statements (see chapter 43), and as arguments to expansion conditions such as **match\_domain**.

Each item in one of these lists is a pattern to be matched against a domain, host, email address, or local part, respectively. In the sections below, the different types of pattern for each case are described, but first we cover some general facilities that apply to all four kinds of list.

Note that other parts of Exim use a *string list* which does not support all the complexity available in domain, host, address and local part lists.

### 10.1 Expansion of lists

Each list is expanded as a single string before it is used.

*Exception: the router headers\_remove option, where list-item splitting is done before string-expansion.*

The result of expansion must be a list, possibly containing empty items, which is split up into separate items for matching. By default, colon is the separator character, but this can be varied if necessary. See sections 6.20 and 6.22 for details of the list syntax; the second of these discusses the way to specify empty list items.

If the string expansion is forced to fail, Exim behaves as if the item it is testing (domain, host, address, or local part) is not in the list. Other expansion failures cause temporary errors.

If an item in a list is a regular expression, backslashes, dollars and possibly other special characters in the expression must be protected against misinterpretation by the string expander. The easiest way to do this is to use the \N expansion feature to indicate that the contents of the regular expression should not be expanded. For example, in an ACL you might have:

```
deny senders = \N^\d{8}\w@.*\baddomain\.example$\N : \
               ${lookup{$domain}lsearch{/badsenders/bydomain}}
```

The first item is a regular expression that is protected from expansion by \N, whereas the second uses the expansion to obtain a list of unwanted senders based on the receiving domain.

### 10.2 Negated items in lists

Items in a list may be positive or negative. Negative items are indicated by a leading exclamation mark, which may be followed by optional white space. A list defines a set of items (domains, etc). When Exim processes one of these lists, it is trying to find out whether a domain, host, address, or local part (respectively) is in the set that is defined by the list. It works like this:

The list is scanned from left to right. If a positive item is matched, the subject that is being checked is in the set; if a negative item is matched, the subject is not in the set. If the end of the list is reached without the subject having matched any of the patterns, it is in the set if the last item was a negative one, but not if it was a positive one. For example, the list in

```
domainlist relay_to_domains = !a.b.c : *.b.c
```

matches any domain ending in *.b.c* except for *a.b.c*. Domains that match neither *a.b.c* nor *\*.b.c* do not match, because the last item in the list is positive. However, if the setting were

```
domainlist relay_to_domains = !a.b.c
```

then all domains other than *a.b.c* would match because the last item in the list is negative. In other words, a list that ends with a negative item behaves as if it had an extra item `: *` on the end.

Another way of thinking about positive and negative items in lists is to read the connector as “or” after a positive item and as “and” after a negative item.

## 10.3 File names in lists

If an item in a domain, host, address, or local part list is an absolute file name (beginning with a slash character), each line of the file is read and processed as if it were an independent item in the list, except that further file names are not allowed, and no expansion of the data from the file takes place. Empty lines in the file are ignored, and the file may also contain comment lines:

- For domain and host lists, if a # character appears anywhere in a line of the file, it and all following characters are ignored.
- Because local parts may legitimately contain # characters, a comment in an address list or local part list file is recognized only if # is preceded by white space or the start of the line. For example:

```
not#comment@x.y.z    # but this is a comment
```

Putting a file name in a list has the same effect as inserting each line of the file as an item in the list (blank lines and comments excepted). However, there is one important difference: the file is read each time the list is processed, so if its contents vary over time, Exim's behaviour changes.

If a file name is preceded by an exclamation mark, the sense of any match within the file is inverted. For example, if

```
hold_domains = !/etc/nohold-domains
```

and the file contains the lines

```
!a.b.c
*.b.c
```

then *a.b.c* is in the set of domains defined by **hold\_domains**, whereas any domain matching *\*.b.c* is not.

## 10.4 An lsearch file is not an out-of-line list

As will be described in the sections that follow, lookups can be used in lists to provide indexed methods of checking list membership. There has been some confusion about the way *lsearch* lookups work in lists. Because an *lsearch* file contains plain text and is scanned sequentially, it is sometimes thought that it is allowed to contain wild cards and other kinds of non-constant pattern. This is not the case. The keys in an *lsearch* file are always fixed strings, just as for any other single-key lookup type.

If you want to use a file to contain wild-card patterns that form part of a list, just give the file name on its own, without a search type, as described in the previous section. You could also use the *wildlsearch* or *nwildlsearch*, but there is no advantage in doing this.

## 10.5 Named lists

A list of domains, hosts, email addresses, or local parts can be given a name which is then used to refer to the list elsewhere in the configuration. This is particularly convenient if the same list is required in several different places. It also allows lists to be given meaningful names, which can improve the readability of the configuration. For example, it is conventional to define a domain list called *local\_domains* for all the domains that are handled locally on a host, using a configuration line such as

```
domainlist local_domains = localhost:my.dom.example
```

Named lists are referenced by giving their name preceded by a plus sign, so, for example, a router that is intended to handle local domains would be configured with the line

```
domains = +local_domains
```

The first router in a configuration is often one that handles all domains except the local ones, using a configuration with a negated item like this:

```
dnslookup:
  driver = dnslookup
  domains = ! +local_domains
```



```
transport = remote_smtp
no_more
```

The four kinds of named list are created by configuration lines starting with the words **domainlist**, **hostlist**, **addresslist**, or **localpartlist**, respectively. Then there follows the name that you are defining, followed by an equals sign and the list itself. For example:

```
hostlist    relay_from_hosts = 192.168.23.0/24 : my.friend.example
addresslist bad_senders = cdb;/etc/badsenders
```

A named list may refer to other named lists:

```
domainlist  dom1 = first.example : second.example
domainlist  dom2 = +dom1 : third.example
domainlist  dom3 = fourth.example : +dom2 : fifth.example
```

**Warning:** If the last item in a referenced list is a negative one, the effect may not be what you intended, because the negation does not propagate out to the higher level. For example, consider:

```
domainlist  dom1 = !a.b
domainlist  dom2 = +dom1 : *.b
```

The second list specifies “either in the **dom1** list or **\*.b**”. The first list specifies just “not *a.b*”, so the domain *x.y* matches it. That means it matches the second list as well. The effect is not the same as

```
domainlist  dom2 = !a.b : *.b
```

where *x.y* does not match. It’s best to avoid negation altogether in referenced lists if you can.

Named lists may have a performance advantage. When Exim is routing an address or checking an incoming message, it caches the result of tests on named lists. So, if you have a setting such as

```
domains = +local_domains
```

on several of your routers or in several ACL statements, the actual test is done only for the first one. However, the caching works only if there are no expansions within the list itself or any sublists that it references. In other words, caching happens only for lists that are known to be the same each time they are referenced.

By default, there may be up to 16 named lists of each type. This limit can be extended by changing a compile-time variable. The use of domain and host lists is recommended for concepts such as local domains, relay domains, and relay hosts. The default configuration is set up like this.

## 10.6 Named lists compared with macros

At first sight, named lists might seem to be no different from macros in the configuration file. However, macros are just textual substitutions. If you write

```
ALIST = host1 : host2
auth_advertise_hosts = !ALIST
```

it probably won’t do what you want, because that is exactly the same as

```
auth_advertise_hosts = !host1 : host2
```

Notice that the second host name is not negated. However, if you use a host list, and write

```
hostlist alist = host1 : host2
auth_advertise_hosts = ! +alist
```

the negation applies to the whole list, and so that is equivalent to

```
auth_advertise_hosts = !host1 : !host2
```

## 10.7 Named list caching

While processing a message, Exim caches the result of checking a named list if it is sure that the list is the same each time. In practice, this means that the cache operates only if the list contains no \$

characters, which guarantees that it will not change when it is expanded. Sometimes, however, you may have an expanded list that you know will be the same each time within a given message. For example:

```
domainlist special_domains = \  
    ${lookup{$sender_host_address}cdb{/some/file}}
```

This provides a list of domains that depends only on the sending host's IP address. If this domain list is referenced a number of times (for example, in several ACL lines, or in several routers) the result of the check is not cached by default, because Exim does not know that it is going to be the same list each time.

By appending `_cache` to `domainlist` you can tell Exim to go ahead and cache the result anyway. For example:

```
domainlist_cache special_domains = ${lookup{...
```

If you do this, you should be absolutely sure that caching is going to do the right thing in all cases. When in doubt, leave it out.

## 10.8 Domain lists

Domain lists contain patterns that are to be matched against a mail domain. The following types of item may appear in domain lists:

- If a pattern consists of a single `@` character, it matches the local host name, as set by the **primary\_hostname** option (or defaulted). This makes it possible to use the same configuration file on several different hosts that differ only in their names.
- If a pattern consists of the string `@[]` it matches an IP address enclosed in square brackets (as in an email address that contains a domain literal), but only if that IP address is recognized as local for email routing purposes. The **local\_interfaces** and **extra\_local\_interfaces** options can be used to control which of a host's several IP addresses are treated as local. In today's Internet, the use of domain literals is controversial.
- If a pattern consists of the string `@mx_any` it matches any domain that has an MX record pointing to the local host or to any host that is listed in **hosts\_treat\_as\_local**. The items `@mx_primary` and `@mx_secondary` are similar, except that the first matches only when a primary MX target is the local host, and the second only when no primary MX target is the local host, but a secondary MX target is. "Primary" means an MX record with the lowest preference value – there may of course be more than one of them.

The MX lookup that takes place when matching a pattern of this type is performed with the resolver options for widening names turned off. Thus, for example, a single-component domain will *not* be expanded by adding the resolver's default domain. See the **qualify\_single** and **search\_parents** options of the *dnslookup* router for a discussion of domain widening.

Sometimes you may want to ignore certain IP addresses when using one of these patterns. You can specify this by following the pattern with `/ignore=<ip list>`, where `<ip list>` is a list of IP addresses. These addresses are ignored when processing the pattern (compare the **ignore\_target\_hosts** option on a router). For example:

```
domains = @mx_any/ignore=127.0.0.1
```

This example matches any domain that has an MX record pointing to one of the local host's IP addresses other than 127.0.0.1.

The list of IP addresses is in fact processed by the same code that processes host lists, so it may contain CIDR-coded network specifications and it may also contain negative items.

Because the list of IP addresses is a sublist within a domain list, you have to be careful about delimiters if there is more than one address. Like any other list, the default delimiter can be changed. Thus, you might have:

```
domains = @mx_any/ignore=<;127.0.0.1;0.0.0.0 : \
an.other.domain : ...
```

so that the sublist uses semicolons for delimiters. When IPv6 addresses are involved, it is easiest to change the delimiter for the main list as well:

```
domains = <? @mx_any/ignore=<;127.0.0.1;::1 ? \
an.other.domain ? ...
```

- If a pattern starts with an asterisk, the remaining characters of the pattern are compared with the terminating characters of the domain. The use of “\*” in domain lists differs from its use in partial matching lookups. In a domain list, the character following the asterisk need not be a dot, whereas partial matching works only in terms of dot-separated components. For example, a domain list item such as *\*key.ex* matches *donkey.ex* as well as *cipherkey.ex*.
- If a pattern starts with a circumflex character, it is treated as a regular expression, and matched against the domain using a regular expression matching function. The circumflex is treated as part of the regular expression. Email domains are case-independent, so this regular expression match is by default case-independent, but you can make it case-dependent by starting it with *(?-i)*. References to descriptions of the syntax of regular expressions are given in chapter 8.

**Warning:** Because domain lists are expanded before being processed, you must escape any backslash and dollar characters in the regular expression, or use the special *\N* sequence (see chapter 11) to specify that it is not to be expanded (unless you really do want to build a regular expression by expansion, of course).

- If a pattern starts with the name of a single-key lookup type followed by a semicolon (for example, “dbm;” or “lsearch;”), the remainder of the pattern must be a file name in a suitable format for the lookup type. For example, for “cdb;” it must be an absolute path:

```
domains = cdb;/etc/mail/local_domains.cdb
```

The appropriate type of lookup is done on the file using the domain name as the key. In most cases, the data that is looked up is not used; Exim is interested only in whether or not the key is present in the file. However, when a lookup is used for the **domains** option on a router or a **domains** condition in an ACL statement, the data is preserved in the *\$domain\_data* variable and can be referred to in other router options or other statements in the same ACL.

- Any of the single-key lookup type names may be preceded by *partial<n>-*, where the *<n>* is optional, for example,

```
domains = partial-dbm;/partial/domains
```

This causes partial matching logic to be invoked; a description of how this works is given in section 9.7.

- Any of the single-key lookup types may be followed by an asterisk. This causes a default lookup for a key consisting of a single asterisk to be done if the original lookup fails. This is not a useful feature when using a domain list to select particular domains (because any domain would match), but it might have value if the result of the lookup is being used via the *\$domain\_data* expansion variable.
- If the pattern starts with the name of a query-style lookup type followed by a semicolon (for example, “nisplus;” or “ldap;”), the remainder of the pattern must be an appropriate query for the lookup type, as described in chapter 9. For example:

```
hold_domains = mysql;select domain from holdlist \
where domain = '${quote_mysql:$domain}';
```

In most cases, the data that is looked up is not used (so for an SQL query, for example, it doesn’t matter what field you select). Exim is interested only in whether or not the query succeeds. However, when a lookup is used for the **domains** option on a router, the data is preserved in the *\$domain\_data* variable and can be referred to in other options.

- If none of the above cases apply, a caseless textual comparison is made between the pattern and the domain.

Here is an example that uses several different kinds of pattern:

```
domainlist funny_domains = \  
  @ : \  
  lib.unseen.edu : \  
  *.foundation.fict.example : \  
  \N^[1-2]\d{3}\.fict\.example$\N : \  
  partial-dbm;/opt/data/penguin/book : \  
  nis;domains.byname : \  
  nisplus;[name=$domain,status=local],domains.org_dir
```

There are obvious processing trade-offs among the various matching modes. Using an asterisk is faster than a regular expression, and listing a few names explicitly probably is too. The use of a file or database lookup is expensive, but may be the only option if hundreds of names are required. Because the patterns are tested in order, it makes sense to put the most commonly matched patterns earlier.

## 10.9 Host lists

Host lists are used to control what remote hosts are allowed to do. For example, some hosts may be allowed to use the local host as a relay, and some may be permitted to use the SMTP ETRN command. Hosts can be identified in two different ways, by name or by IP address. In a host list, some types of pattern are matched to a host name, and some are matched to an IP address. You need to be particularly careful with this when single-key lookups are involved, to ensure that the right value is being used as the key.

### 10.10 Special host list patterns

If a host list item is the empty string, it matches only when no remote host is involved. This is the case when a message is being received from a local process using SMTP on the standard input, that is, when a TCP/IP connection is not used.

The special pattern “\*” in a host list matches any host or no host. Neither the IP address nor the name is actually inspected.

### 10.11 Host list patterns that match by IP address

If an IPv4 host calls an IPv6 host and the call is accepted on an IPv6 socket, the incoming address actually appears in the IPv6 host as `::ffff:<v4address>`. When such an address is tested against a host list, it is converted into a traditional IPv4 address first. (Not all operating systems accept IPv4 calls on IPv6 sockets, as there have been some security concerns.)

The following types of pattern in a host list check the remote host by inspecting its IP address:

- If the pattern is a plain domain name (not a regular expression, not starting with \*, not a lookup of any kind), Exim calls the operating system function to find the associated IP address(es). Exim uses the newer *getipnodebyname()* function when available, otherwise *gethostbyname()*. This typically causes a forward DNS lookup of the name. The result is compared with the IP address of the subject host.

If there is a temporary problem (such as a DNS timeout) with the host name lookup, a temporary error occurs. For example, if the list is being used in an ACL condition, the ACL gives a “defer” response, usually leading to a temporary SMTP error code. If no IP address can be found for the host name, what happens is described in section 10.14 below.

- If the pattern is “@”, the primary host name is substituted and used as a domain name, as just described.
- If the pattern is an IP address, it is matched against the IP address of the subject host. IPv4 addresses are given in the normal “dotted-quad” notation. IPv6 addresses can be given in colon-separated format, but the colons have to be doubled so as not to be taken as item separators when the default list separator is used. IPv6 addresses are recognized even when Exim is compiled

without IPv6 support. This means that if they appear in a host list on an IPv4-only host, Exim will not treat them as host names. They are just addresses that can never match a client host.

- If the pattern is “@[]”, it matches the IP address of any IP interface on the local host. For example, if the local host is an IPv4 host with one interface address 10.45.23.56, these two ACL statements have the same effect:

```
accept hosts = 127.0.0.1 : 10.45.23.56
accept hosts = @[]
```

- If the pattern is an IP address followed by a slash and a mask length (for example 10.11.42.0/24), it is matched against the IP address of the subject host under the given mask. This allows, an entire network of hosts to be included (or excluded) by a single item. The mask uses CIDR notation; it specifies the number of address bits that must match, starting from the most significant end of the address.

**Note:** The mask is *not* a count of addresses, nor is it the high number of a range of addresses. It is the number of bits in the network portion of the address. The above example specifies a 24-bit netmask, so it matches all 256 addresses in the 10.11.42.0 network. An item such as

```
192.168.23.236/31
```

matches just two addresses, 192.168.23.236 and 192.168.23.237. A mask value of 32 for an IPv4 address is the same as no mask at all; just a single address matches.

Here is another example which shows an IPv4 and an IPv6 network:

```
recipient_unqualified_hosts = 192.168.0.0/16: \
                              3ffe::ffff:836f:::/48
```

The doubling of list separator characters applies only when these items appear inline in a host list. It is not required when indirecting via a file. For example:

```
recipient_unqualified_hosts = /opt/exim/unqualnets
```

could make use of a file containing

```
172.16.0.0/12
3ffe:ffff:836f::/48
```

to have exactly the same effect as the previous example. When listing IPv6 addresses inline, it is usually more convenient to use the facility for changing separator characters. This list contains the same two networks:

```
recipient_unqualified_hosts = <; 172.16.0.0/12; \
                              3ffe:ffff:836f::/48
```

The separator is changed to semicolon by the leading “<;” at the start of the list.

## 10.12 Host list patterns for single-key lookups by host address

When a host is to be identified by a single-key lookup of its complete IP address, the pattern takes this form:

```
net-<single-key-search-type>; <search-data>
```

For example:

```
hosts_lookup = net-cdb;/hosts-by-ip.db
```

The text form of the IP address of the subject host is used as the lookup key. IPv6 addresses are converted to an unabbreviated form, using lower case letters, with dots as separators because colon is the key terminator in *lsearch* files. [Colons can in fact be used in keys in *lsearch* files by quoting the keys, but this is a facility that was added later.] The data returned by the lookup is not used.

Single-key lookups can also be performed using masked IP addresses, using patterns of this form:

```
net<number>-<single-key-search-type>; <search-data>
```

For example:

```
net24-dbm; /networks.db
```

The IP address of the subject host is masked using `<number>` as the mask length. A textual string is constructed from the masked value, followed by the mask, and this is used as the lookup key. For example, if the host's IP address is 192.168.34.6, the key that is looked up for the above example is "192.168.34.0/24".

When an IPv6 address is converted to a string, dots are normally used instead of colons, so that keys in *lsearch* files need not contain colons (which terminate *lsearch* keys). This was implemented some time before the ability to quote keys was made available in *lsearch* files. However, the more recently implemented *iplsearch* files do require colons in IPv6 keys (notated using the quoting facility) so as to distinguish them from IPv4 keys. For this reason, when the lookup type is *iplsearch*, IPv6 addresses are converted using colons and not dots. In all cases, full, unabbreviated IPv6 addresses are always used.

Ideally, it would be nice to tidy up this anomalous situation by changing to colons in all cases, given that quoting is now available for *lsearch*. However, this would be an incompatible change that might break some existing configurations.

**Warning:** Specifying **net32-** (for an IPv4 address) or **net128-** (for an IPv6 address) is not the same as specifying just **net-** without a number. In the former case the key strings include the mask value, whereas in the latter case the IP address is used on its own.

## 10.13 Host list patterns that match by host name

There are several types of pattern that require Exim to know the name of the remote host. These are either wildcard patterns or lookups by name. (If a complete hostname is given without any wildcarding, it is used to find an IP address to match against, as described in section 10.11 above.)

If the remote host name is not already known when Exim encounters one of these patterns, it has to be found from the IP address. Although many sites on the Internet are conscientious about maintaining reverse DNS data for their hosts, there are also many that do not do this. Consequently, a name cannot always be found, and this may lead to unwanted effects. Take care when configuring host lists with wildcarded name patterns. Consider what will happen if a name cannot be found.

Because of the problems of determining host names from IP addresses, matching against host names is not as common as matching against IP addresses.

By default, in order to find a host name, Exim first does a reverse DNS lookup; if no name is found in the DNS, the system function (*gethostbyaddr()* or *getipnodebyaddr()* if available) is tried. The order in which these lookups are done can be changed by setting the **host\_lookup\_order** option. For security, once Exim has found one or more names, it looks up the IP addresses for these names and compares them with the IP address that it started with. Only those names whose IP addresses match are accepted. Any other names are discarded. If no names are left, Exim behaves as if the host name cannot be found. In the most common case there is only one name and one IP address.

There are some options that control what happens if a host name cannot be found. These are described in section 10.14 below.

As a result of aliasing, hosts may have more than one name. When processing any of the following types of pattern, all the host's names are checked:

- If a pattern starts with "\*" the remainder of the item must match the end of the host name. For example, \*.b.c matches all hosts whose names end in .b.c. This special simple form is provided because this is a very common requirement. Other kinds of wildcarding require the use of a regular expression.
- If the item starts with "^" it is taken to be a regular expression which is matched against the host name. Host names are case-independent, so this regular expression match is by default case-independent, but you can make it case-dependent by starting it with (?-i). References to descriptions of the syntax of regular expressions are given in chapter 8. For example,

```
^(a|b)\.c\.d$
```

is a regular expression that matches either of the two hosts *a.c.d* or *b.c.d*. When a regular expression is used in a host list, you must take care that backslash and dollar characters are not misinterpreted as part of the string expansion. The simplest way to do this is to use `\N` to mark that part of the string as non-expandable. For example:

```
sender_unqualified_hosts = \N^(a|b)\.c\.d$\N : ....
```

**Warning:** If you want to match a complete host name, you must include the `$` terminating metacharacter in the regular expression, as in the above example. Without it, a match at the start of the host name is all that is required.

## 10.14 Behaviour when an IP address or name cannot be found

While processing a host list, Exim may need to look up an IP address from a name (see section 10.11), or it may need to look up a host name from an IP address (see section 10.13). In either case, the behaviour when it fails to find the information it is seeking is the same.

**Note:** This section applies to permanent lookup failures. It does *not* apply to temporary DNS errors, whose handling is described in the next section.

Exim parses a host list from left to right. If it encounters a permanent lookup failure in any item in the host list before it has found a match, Exim treats it as a failure and the default behavior is as if the host does not match the list. This may not always be what you want to happen. To change Exim's behaviour, the special items `+include_unknown` or `+ignore_unknown` may appear in the list (at top level – they are not recognized in an indirected file).

- If any item that follows `+include_unknown` requires information that cannot be found, Exim behaves as if the host does match the list. For example,

```
host_reject_connection = +include_unknown:*.enemy.ex
```

rejects connections from any host whose name matches `*.enemy.ex`, and also any hosts whose name it cannot find.

- If any item that follows `+ignore_unknown` requires information that cannot be found, Exim ignores that item and proceeds to the rest of the list. For example:

```
accept hosts = +ignore_unknown : friend.example : \
192.168.4.5
```

accepts from any host whose name is *friend.example* and from 192.168.4.5, whether or not its host name can be found. Without `+ignore_unknown`, if no name can be found for 192.168.4.5, it is rejected.

Both `+include_unknown` and `+ignore_unknown` may appear in the same list. The effect of each one lasts until the next, or until the end of the list.

## 10.15 Mixing wildcarded host names and addresses in host lists

This section explains the host/ip processing logic with the same concepts as the previous section, but specifically addresses what happens when a wildcarded hostname is one of the items in the hostlist.

- If you have name lookups or wildcarded host names and IP addresses in the same host list, you should normally put the IP addresses first. For example, in an ACL you could have:

```
accept hosts = 10.9.8.7 : *.friend.example
```

The reason you normally would order it this way lies in the left-to-right way that Exim processes lists. It can test IP addresses without doing any DNS lookups, but when it reaches an item that requires a host name, it fails if it cannot find a host name to compare with the pattern. If the above list is given in the opposite order, the **accept** statement fails for a host whose name cannot be found, even if its IP address is 10.9.8.7.

- If you really do want to do the name check first, and still recognize the IP address, you can rewrite the ACL like this:

```
accept hosts = *.friend.example
accept hosts = 10.9.8.7
```

If the first **accept** fails, Exim goes on to try the second one. See chapter 43 for details of ACLs. Alternatively, you can use `+ignore_unknown`, which was discussed in depth in the first example in this section.

## 10.16 Temporary DNS errors when looking up host information

A temporary DNS lookup failure normally causes a defer action (except when **dns\_again\_means\_nonexist** converts it into a permanent error). However, host lists can include `+ignore_defer` and `+include_defer`, analogous to `+ignore_unknown` and `+include_unknown`, as described in the previous section. These options should be used with care, probably only in non-critical host lists such as whitelists.

## 10.17 Host list patterns for single-key lookups by host name

If a pattern is of the form

```
<single-key-search-type>;<search-data>
```

for example

```
dbm;/host/accept/list
```

a single-key lookup is performed, using the host name as its key. If the lookup succeeds, the host matches the item. The actual data that is looked up is not used.

**Reminder:** With this kind of pattern, you must have host *names* as keys in the file, not IP addresses. If you want to do lookups based on IP addresses, you must precede the search type with “net-” (see section 10.12). There is, however, no reason why you could not use two items in the same list, one doing an address lookup and one doing a name lookup, both using the same file.

## 10.18 Host list patterns for query-style lookups

If a pattern is of the form

```
<query-style-search-type>;<query>
```

the query is obeyed, and if it succeeds, the host matches the item. The actual data that is looked up is not used. The variables `$sender_host_address` and `$sender_host_name` can be used in the query. For example:

```
hosts_lookup = pgsql;\
select ip from hostlist where ip='$sender_host_address'
```

The value of `$sender_host_address` for an IPv6 address contains colons. You can use the **sg** expansion item to change this if you need to. If you want to use masked IP addresses in database queries, you can use the **mask** expansion operator.

If the query contains a reference to `$sender_host_name`, Exim automatically looks up the host name if it has not already done so. (See section 10.13 for comments on finding host names.)

Historical note: prior to release 4.30, Exim would always attempt to find a host name before running the query, unless the search type was preceded by `net-`. This is no longer the case. For backwards compatibility, `net-` is still recognized for query-style lookups, but its presence or absence has no effect. (Of course, for single-key lookups, `net-` is important. See section 10.12.)



## 10.19 Address lists

Address lists contain patterns that are matched against mail addresses. There is one special case to be considered: the sender address of a bounce message is always empty. You can test for this by providing an empty item in an address list. For example, you can set up a router to process bounce messages by using this option setting:

```
senders = :
```

The presence of the colon creates an empty item. If you do not provide any data, the list is empty and matches nothing. The empty sender can also be detected by a regular expression that matches an empty string, and by a query-style lookup that succeeds when *\$sender\_address* is empty.

Non-empty items in an address list can be straightforward email addresses. For example:

```
senders = jbc@askone.example : hs@anacreon.example
```

A certain amount of wildcarding is permitted. If a pattern contains an @ character, but is not a regular expression and does not begin with a semicolon-terminated lookup type (described below), the local part of the subject address is compared with the local part of the pattern, which may start with an asterisk. If the local parts match, the domain is checked in exactly the same way as for a pattern in a domain list. For example, the domain can be wildcarded, refer to a named list, or be a lookup:

```
deny senders = *@*.spamming.site:\
               *@+hostile_domains:\
               bozo@partial-lsearch;/list/of/dodgy/sites:\
               *@dbm;/bad/domains.db
```

If a local part that begins with an exclamation mark is required, it has to be specified using a regular expression, because otherwise the exclamation mark is treated as a sign of negation, as is standard in lists.

If a non-empty pattern that is not a regular expression or a lookup does not contain an @ character, it is matched against the domain part of the subject address. The only two formats that are recognized this way are a literal domain, or a domain pattern that starts with \*. In both these cases, the effect is the same as if \*@ preceded the pattern. For example:

```
deny senders = enemy.domain : *.enemy.domain
```

The following kinds of more complicated address list pattern can match any address, including the empty address that is characteristic of bounce message senders:

- If (after expansion) a pattern starts with “^”, a regular expression match is done against the complete address, with the pattern as the regular expression. You must take care that backslash and dollar characters are not misinterpreted as part of the string expansion. The simplest way to do this is to use \N to mark that part of the string as non-expandable. For example:

```
deny senders = \N^.*this.*@example\.com$\N : \
               \N^d{8}.+@spamhaus.example$\N : ...
```

The \N sequences are removed by the expansion, so these items do indeed start with “^” by the time they are being interpreted as address patterns.

- Complete addresses can be looked up by using a pattern that starts with a lookup type terminated by a semicolon, followed by the data for the lookup. For example:

```
deny senders = cdb;/etc/blocked.senders : \
               mysql;select address from blocked where \
               address='${quote_mysql:$sender_address}'
```

Both query-style and single-key lookup types can be used. For a single-key lookup type, Exim uses the complete address as the key. However, empty keys are not supported for single-key lookups, so a match against the empty address always fails. This restriction does not apply to query-style lookups.

Partial matching for single-key lookups (section 9.7) cannot be used, and is ignored if specified, with an entry being written to the panic log. However, you can configure lookup defaults, as described in section 9.6, but this is useful only for the “\*@” type of default. For example, with this lookup:

```
accept senders = lsearch*@;/some/file
```

the file could contain lines like this:

```
user1@domain1.example
*@domain2.example
```

and for the sender address *nimrod@jaeger.example*, the sequence of keys that are tried is:

```
nimrod@jaeger.example
*@jaeger.example
*
```

**Warning 1:** Do not include a line keyed by “\*” in the file, because that would mean that every address matches, thus rendering the test useless.

**Warning 2:** Do not confuse these two kinds of item:

```
deny recipients = dbm*@;/some/file
deny recipients = *@dbm;/some/file
```

The first does a whole address lookup, with defaulting, as just described, because it starts with a lookup type. The second matches the local part and domain independently, as described in a bullet point below.

The following kinds of address list pattern can match only non-empty addresses. If the subject address is empty, a match against any of these pattern types always fails.

- If a pattern starts with “@@” followed by a single-key lookup item (for example, @@lsearch;/some/file), the address that is being checked is split into a local part and a domain. The domain is looked up in the file. If it is not found, there is no match. If it is found, the data that is looked up from the file is treated as a colon-separated list of local part patterns, each of which is matched against the subject local part in turn.

The lookup may be a partial one, and/or one involving a search for a default keyed by “\*” (see section 9.6). The local part patterns that are looked up can be regular expressions or begin with “\*”, or even be further lookups. They may also be independently negated. For example, with

```
deny senders = @@dbm;/etc/reject-by-domain
```

the data from which the DBM file is built could contain lines like

```
baddomain.com: !postmaster : *
```

to reject all senders except **postmaster** from that domain.

If a local part that actually begins with an exclamation mark is required, it has to be specified using a regular expression. In *lsearch* files, an entry may be split over several lines by indenting the second and subsequent lines, but the separating colon must still be included at line breaks. White space surrounding the colons is ignored. For example:

```
aol.com: spammer1 : spammer2 : ^[0-9]+$ :
      spammer3 : spammer4
```

As in all colon-separated lists in Exim, a colon can be included in an item by doubling.

If the last item in the list starts with a right angle-bracket, the remainder of the item is taken as a new key to look up in order to obtain a continuation list of local parts. The new key can be any sequence of characters. Thus one might have entries like

```
aol.com: spammer1 : spammer 2 : >*
xyz.com: spammer3 : >*
*: ^\d{8}$
```

in a file that was searched with `@@dbm*`, to specify a match for 8-digit local parts for all domains, in addition to the specific local parts listed for each domain. Of course, using this feature costs another lookup each time a chain is followed, but the effort needed to maintain the data is reduced.

It is possible to construct loops using this facility, and in order to catch them, the chains may be no more than fifty items long.

- The `@@<lookup>` style of item can also be used with a query-style lookup, but in this case, the chaining facility is not available. The lookup can only return a single list of local parts.

**Warning:** There is an important difference between the address list items in these two examples:

```
senders = +my_list
senders = *@+my_list
```

In the first one, `my_list` is a named address list, whereas in the second example it is a named domain list.

## 10.20 Case of letters in address lists

Domains in email addresses are always handled caselessly, but for local parts case may be significant on some systems (see **caseful\_local\_part** for how Exim deals with this when routing addresses). However, RFC 2505 (*Anti-Spam Recommendations for SMTP MTAs*) suggests that matching of addresses to blocking lists should be done in a case-independent manner. Since most address lists in Exim are used for this kind of control, Exim attempts to do this by default.

The domain portion of an address is always lowercased before matching it to an address list. The local part is lowercased by default, and any string comparisons that take place are done caselessly. This means that the data in the address list itself, in files included as plain file names, and in any file that is looked up using the “@@” mechanism, can be in any case. However, the keys in files that are looked up by a search type other than *lsearch* (which works caselessly) must be in lower case, because these lookups are not case-independent.

To allow for the possibility of caseful address list matching, if an item in an address list is the string “+caseful”, the original case of the local part is restored for any comparisons that follow, and string comparisons are no longer case-independent. This does not affect the domain, which remains in lower case. However, although independent matches on the domain alone are still performed caselessly, regular expressions that match against an entire address become case-sensitive after “+caseful” has been seen.

## 10.21 Local part lists

Case-sensitivity in local part lists is handled in the same way as for address lists, as just described. The “+caseful” item can be used if required. In a setting of the **local\_parts** option in a router with **caseful\_local\_part** set false, the subject is lowercased and the matching is initially case-insensitive. In this case, “+caseful” will restore case-sensitive matching in the local part list, but not elsewhere in the router. If **caseful\_local\_part** is set true in a router, matching in the **local\_parts** option is case-sensitive from the start.

If a local part list is indirected to a file (see section 10.3), comments are handled in the same way as address lists – they are recognized only if the # is preceded by white space or the start of the line. Otherwise, local part lists are matched in the same way as domain lists, except that the special items that refer to the local host (`@`, `@[]`, `@mx_any`, `@mx_primary`, and `@mx_secondary`) are not recognized. Refer to section 10.8 for details of the other available item types.

# 11. String expansions

Many strings in Exim's run time configuration are expanded before use. Some of them are expanded every time they are used; others are expanded only once.

When a string is being expanded it is copied verbatim from left to right except when a dollar or backslash character is encountered. A dollar specifies the start of a portion of the string that is interpreted and replaced as described below in section 11.5 onwards. Backslash is used as an escape character, as described in the following section.

Whether a string is expanded depends upon the context. Usually this is solely dependent upon the option for which a value is sought; in this documentation, options for which string expansion is performed are marked with † after the data type. ACL rules always expand strings. A couple of expansion conditions do not expand some of the brace-delimited branches, for security reasons.

## 11.1 Literal text in expanded strings

An uninterpreted dollar can be included in an expanded string by putting a backslash in front of it. A backslash can be used to prevent any special character being treated specially in an expansion, including backslash itself. If the string appears in quotes in the configuration file, two backslashes are required because the quotes themselves cause interpretation of backslashes when the string is read in (see section 6.17).

A portion of the string can be specified as non-expandable by placing it between two occurrences of \N. This is particularly useful for protecting regular expressions, which often contain backslashes and dollar signs. For example:

```
deny senders = \N^\d{8}[a-z]@some\.site\.example$\N
```

On encountering the first \N, the expander copies subsequent characters without interpretation until it reaches the next \N or the end of the string.

## 11.2 Character escape sequences in expanded strings

A backslash followed by one of the letters “n”, “r”, or “t” in an expanded string is recognized as an escape sequence for the character newline, carriage return, or tab, respectively. A backslash followed by up to three octal digits is recognized as an octal encoding for a single character, and a backslash followed by “x” and up to two hexadecimal digits is a hexadecimal encoding.

These escape sequences are also recognized in quoted strings when they are read in. Their interpretation in expansions as well is useful for unquoted strings, and for other cases such as looked-up strings that are then expanded.

## 11.3 Testing string expansions

Many expansions can be tested by calling Exim with the **-be** option. This takes the command arguments, or lines from the standard input if there are no arguments, runs them through the string expansion code, and writes the results to the standard output. Variables based on configuration values are set up, but since no message is being processed, variables such as *\$local\_part* have no value. Nevertheless the **-be** option can be useful for checking out file and database lookups, and the use of expansion operators such as **sg**, **substr** and **nhash**.

Exim gives up its root privilege when it is called with the **-be** option, and instead runs under the uid and gid it was called with, to prevent users from using **-be** for reading files to which they do not have access.

If you want to test expansions that include variables whose values are taken from a message, there are two other options that can be used. The **-bem** option is like **-be** except that it is followed by a file name. The file is read as a message before doing the test expansions. For example:

```
exim -bem /tmp/test.message '$h_subject:'
```

The **-Mset** option is used in conjunction with **-be** and is followed by an Exim message identifier. For example:

```
exim -be -Mset 1GrA8W-0004WS-LQ '$recipients'
```

This loads the message from Exim's spool before doing the test expansions, and is therefore restricted to admin users.

## 11.4 Forced expansion failure

A number of expansions that are described in the following section have alternative “true” and “false” substrings, enclosed in brace characters (which are sometimes called “curly brackets”). Which of the two strings is used depends on some condition that is evaluated as part of the expansion. If, instead of a “false” substring, the word “fail” is used (not in braces), the entire string expansion fails in a way that can be detected by the code that requested the expansion. This is called “forced expansion failure”, and its consequences depend on the circumstances. In some cases it is no different from any other expansion failure, but in others a different action may be taken. Such variations are mentioned in the documentation of the option that is being expanded.

## 11.5 Expansion items

The following items are recognized in expanded strings. White space may be used between sub-items that are keywords or substrings enclosed in braces inside an outer set of braces, to improve readability. **Warning:** Within braces, white space is significant.

**`$<variable name>`** or **`${<variable name>}`**

Substitute the contents of the named variable, for example:

```
$local_part  
${domain}
```

The second form can be used to separate the name from subsequent alphanumeric characters. This form (using braces) is available only for variables; it does *not* apply to message headers. The names of the variables are given in section 11.9 below. If the name of a non-existent variable is given, the expansion fails.

**`${<op>:<string>}`**

The string is first itself expanded, and then the operation specified by `<op>` is applied to it. For example:

```
${lc:$local_part}
```

The string starts with the first character after the colon, which may be leading white space. A list of operators is given in section 11.6 below. The operator notation is used for simple expansion items that have just one argument, because it reduces the number of braces and therefore makes the string easier to understand.

**`$bheader_<header name>:`** or **`$bh_<header name>:`**

This item inserts “basic” header lines. It is described with the **header** expansion item below.

**`${acl{<name>}{<arg>}...}`**

The name and zero to nine argument strings are first expanded separately. The expanded arguments are assigned to the variables `$acl_arg1` to `$acl_arg9` in order. Any unused are made empty. The variable `$acl_narg` is set to the number of arguments. The named ACL (see chapter 43) is called and may use the variables; if another acl expansion is used the values are restored after it returns. If the ACL sets a value using a “message =” modifier and returns accept or deny, the value becomes the result of the expansion. If no message is set and the ACL returns accept or deny the expansion result is an empty string. If the ACL returns defer the result is a forced-fail. Otherwise the expansion fails.

**`${certextract{<field>}{<certificate>}{<string2>}{<string3>}}`**

The `<certificate>` must be a variable of type certificate. The field name is expanded and used to retrieve the relevant field from the certificate. Supported fields are:

version	
serial_number	
subject	RFC4514 DN
issuer	RFC4514 DN
notbefore	time
notafter	time
sig_algorithm	
signature	
subj_altname	tagged list
ocsp_uri	list
crl_uri	list

If the field is found, `<string2>` is expanded, and replaces the whole item; otherwise `<string3>` is used. During the expansion of `<string2>` the variable `$value` contains the value that has been extracted. Afterwards, it is restored to any previous value it might have had.

If `{<string3>}` is omitted, the item is replaced by an empty string if the key is not found. If `{<string2>}` is also omitted, the value that was extracted is used.

Some field names take optional modifiers, appended and separated by commas.

The field selectors marked as "RFC4514" above output a Distinguished Name string which is not quite parseable by Exim as a comma-separated tagged list (the exceptions being elements containing commas). RDN elements of a single type may be selected by a modifier of the type label; if so the expansion result is a list (newline-separated by default). The separator may be changed by another modifier of a right angle-bracket followed immediately by the new separator. Recognised RDN type labels include "CN", "O", "OU" and "DC".

The field selectors marked as "time" above take an optional modifier of "int" for which the result is the number of seconds since epoch. Otherwise the result is a human-readable string in the timezone selected by the main "timezone" option.

The field selectors marked as "list" above return a list, newline-separated by default, (embedded separator characters in elements are doubled). The separator may be changed by a modifier of a right angle-bracket followed immediately by the new separator.

The field selectors marked as "tagged" above prefix each list element with a type string and an equals sign. Elements of only one type may be selected by a modifier which is one of "dns", "uri" or "mail"; if so the element tags are omitted.

If not otherwise noted field values are presented in human-readable form.

**`$(dlfunc{<file>}{<function>}{<arg>}{<arg>}...}`**

This expansion dynamically loads and then calls a locally-written C function. This functionality is available only if Exim is compiled with

```
EXPAND_DLFUNC=yes
```

set in *Local/Makefile*. Once loaded, Exim remembers the dynamically loaded object so that it doesn't reload the same object file in the same Exim process (but of course Exim does start new processes frequently).

There may be from zero to eight arguments to the function. When compiling a local function that is to be called in this way, *local\_scan.h* should be included. The Exim variables and functions that are defined by that API are also available for dynamically loaded functions. The function itself must have the following type:

```
int dlfunction(uschar **yield, int argc, uschar *argv[])
```

Where `uschar` is a typedef for `unsigned char` in *local\_scan.h*. The function should return one of the following values:

OK: Success. The string that is placed in the variable `yield` is put into the expanded string that is being built.

FAIL: A non-forced expansion failure occurs, with the error message taken from `yield`, if it is set.

FAIL\_FORCED: A forced expansion failure occurs, with the error message taken from *yield* if it is set.

ERROR: Same as FAIL, except that a panic log entry is written.

When compiling a function that is to be used in this way with gcc, you need to add **-shared** to the gcc command. Also, in the Exim build-time configuration, you must add **-export-dynamic** to EXTRALIBS.

**`${env{<key>}{<string1>}{<string2>}}`**

The key is first expanded separately, and leading and trailing white space removed. This is then searched for as a name in the environment. If a variable is found then its value is placed in *\$value* and *<string1>* is expanded, otherwise *<string2>* is expanded.

Instead of *{<string2>}* the word “fail” (not in curly brackets) can appear, for example:

```
${env{USER}{$value} fail }
```

This forces an expansion failure (see section 11.4); *{<string1>}* must be present for “fail” to be recognized.

If *{<string2>}* is omitted an empty string is substituted on search failure. If *{<string1>}* is omitted the search result is substituted on search success.

The environment is adjusted by the **keep\_environment** and **add\_environment** main section options.

**`${extract{<key>}{<string1>}{<string2>}{<string3>}}`**

The key and *<string1>* are first expanded separately. Leading and trailing white space is removed from the key (but not from any of the strings). The key must not be empty and must not consist entirely of digits. The expanded *<string1>* must be of the form:

```
<key1> = <value1> <key2> = <value2> ...
```

where the equals signs and spaces (but not both) are optional. If any of the values contain white space, they must be enclosed in double quotes, and any values that are enclosed in double quotes are subject to escape processing as described in section 6.17. The expanded *<string1>* is searched for the value that corresponds to the key. The search is case-insensitive. If the key is found, *<string2>* is expanded, and replaces the whole item; otherwise *<string3>* is used. During the expansion of *<string2>* the variable *\$value* contains the value that has been extracted. Afterwards, it is restored to any previous value it might have had.

If *{<string3>}* is omitted, the item is replaced by an empty string if the key is not found. If *{<string2>}* is also omitted, the value that was extracted is used. Thus, for example, these two expansions are identical, and yield “2001”:

```
${extract{gid}{uid=1984 gid=2001}}  
${extract{gid}{uid=1984 gid=2001}{$value}}
```

Instead of *{<string3>}* the word “fail” (not in curly brackets) can appear, for example:

```
${extract{Z}{A=... B=...}{$value} fail }
```

This forces an expansion failure (see section 11.4); *{<string2>}* must be present for “fail” to be recognized.

**`${extract{<number>}{<separators>}{<string1>}{<string2>}{<string3>}}`**

The *<number>* argument must consist entirely of decimal digits, apart from leading and trailing white space, which is ignored. This is what distinguishes this form of **extract** from the previous kind. It behaves in the same way, except that, instead of extracting a named field, it extracts from *<string1>* the field whose number is given as the first argument. You can use *\$value* in *<string2>* or *fail* instead of *<string3>* as before.

The fields in the string are separated by any one of the characters in the separator string. These may include space or tab characters. The first field is numbered one. If the number is negative, the fields are counted from the end of the string, with the rightmost one numbered -1. If the number given is zero, the entire string is returned. If the modulus of the number is greater than the number

of fields in the string, the result is the expansion of `<string3>`, or the empty string if `<string3>` is not provided. For example:

```
{extract{2}{:}{x:42:99:& Mailer::/bin/bash}}
```

yields “42”, and

```
{extract{-4}{:}{x:42:99:& Mailer::/bin/bash}}
```

yields “99”. Two successive separators mean that the field between them is empty (for example, the fifth field above).

### **`${filter{<string>}{<condition>}}`**

After expansion, `<string>` is interpreted as a list, colon-separated by default, but the separator can be changed in the usual way. For each item in this list, its value is place in `$item`, and then the condition is evaluated. If the condition is true, `$item` is added to the output as an item in a new list; if the condition is false, the item is discarded. The separator used for the output list is the same as the one used for the input, but a separator setting is not included in the output. For example:

```
{filter{a:b:c}{!eq{$item}{b}}
```

yields `a:c`. At the end of the expansion, the value of `$item` is restored to what it was before. See also the **map** and **reduce** expansion items.

### **`${hash{<string1>}{<string2>}{<string3>}}`**

This is a textual hashing function, and was the first to be implemented in early versions of Exim. In current releases, there are other hashing functions (numeric, MD5, and SHA-1), which are described below.

The first two strings, after expansion, must be numbers. Call them `<m>` and `<n>`. If you are using fixed values for these numbers, that is, if `<string1>` and `<string2>` do not change when they are expanded, you can use the simpler operator notation that avoids some of the braces:

```
{hash_<n>_<m>:<string>}
```

The second number is optional (in both notations). If `<n>` is greater than or equal to the length of the string, the expansion item returns the string. Otherwise it computes a new string of length `<n>` by applying a hashing function to the string. The new string consists of characters taken from the first `<m>` characters of the string

```
abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789
```

If `<m>` is not present the value 26 is used, so that only lower case letters appear. For example:

```
{hash{3}{monty}}           yields jmg
{hash{5}{monty}}           yields monty
{hash{4}{62}{monty python}} yields fbWx
```

### **`$header_<header name>:` or **`$h_<header name>:`****

### **`$bheader_<header name>:` or **`$bh_<header name>:`****

### **`$rheader_<header name>:` or **`$rh_<header name>:`****

Substitute the contents of the named message header line, for example

```
$header_reply-to:
```

The newline that terminates a header line is not included in the expansion, but internal newlines (caused by splitting the header line over several physical lines) may be present.

The difference between **rheader**, **bheader**, and **header** is in the way the data in the header line is interpreted.

- **rheader** gives the original “raw” content of the header line, with no processing at all, and without the removal of leading and trailing white space.
- **bheader** removes leading and trailing white space, and then decodes base64 or quoted-printable MIME “words” within the header text, but does no character set translation. If decoding of what looks superficially like a MIME “word” fails, the raw string is returned. If decoding produces a



binary zero character, it is replaced by a question mark – this is what Exim does for binary zeros that are actually received in header lines.

- **header** tries to translate the string as decoded by **bheader** to a standard character set. This is an attempt to produce the same string as would be displayed on a user's MUA. If translation fails, the **bheader** string is returned. Translation is attempted only on operating systems that support the *iconv()* function. This is indicated by the compile-time macro HAVE\_ICONV in a system Makefile or in *Local/Makefile*.

In a filter file, the target character set for **header** can be specified by a command of the following form:

```
headers charset "UTF-8"
```

This command affects all references to *\$h\_* (or *\$header\_*) expansions in subsequently obeyed filter commands. In the absence of this command, the target character set in a filter is taken from the setting of the **headers\_charset** option in the runtime configuration. The value of this option defaults to the value of HEADERS\_CHARSET in *Local/Makefile*. The ultimate default is ISO-8859-1.

Header names follow the syntax of RFC 2822, which states that they may contain any printing characters except space and colon. Consequently, curly brackets *do not* terminate header names, and should not be used to enclose them as if they were variables. Attempting to do so causes a syntax error.

Only header lines that are common to all copies of a message are visible to this mechanism. These are the original header lines that are received with the message, and any that are added by an ACL statement or by a system filter. Header lines that are added to a particular copy of a message by a router or transport are not accessible.

For incoming SMTP messages, no header lines are visible in ACLs that are obeyed before the data phase completes, because the header structure is not set up until the message is received. They are visible in DKIM, PRDR and DATA ACLs. Header lines that are added in a RCPT ACL (for example) are saved until the message's incoming header lines are available, at which point they are added. When any of the above ACLs are running, however, header lines added by earlier ACLs are visible.

Upper case and lower case letters are synonymous in header names. If the following character is white space, the terminating colon may be omitted, but this is not recommended, because you may then forget it when it is needed. When white space terminates the header name, this white space is included in the expanded string. If the message does not contain the given header, the expansion item is replaced by an empty string. (See the **def** condition in section 11.7 for a means of testing for the existence of a header.)

If there is more than one header with the same name, they are all concatenated to form the substitution string, up to a maximum length of 64K. Unless **rheader** is being used, leading and trailing white space is removed from each header before concatenation, and a completely empty header is ignored. A newline character is then inserted between non-empty headers, but there is no newline at the very end. For the **header** and **bheader** expansion, for those headers that contain lists of addresses, a comma is also inserted at the junctions between headers. This does not happen for the **rheader** expansion.

**`${hmac{<hashname>}{<secret>}{<string>}}`**

This function uses cryptographic hashing (either MD5 or SHA-1) to convert a shared secret and some text into a message authentication code, as specified in RFC 2104. This differs from `${md5:secret_text...}` or `${sha1:secret_text...}` in that the hmac step adds a signature to the cryptographic hash, allowing for authentication that is not possible with MD5 or SHA-1 alone. The hash name must expand to either `md5` or `sha1` at present. For example:

```
${hmac{md5}{somesecret}{$primary_hostname $tod_log}}
```

For the hostname *mail.example.com* and time 2002-10-17 11:30:59, this produces:

```
dd97e3ba5d1a61b5006108f8c8252953
```

As an example of how this might be used, you might put in the main part of an Exim configuration:

```
SPAMSCAN_SECRET=coghheeLei2thahw
```

In a router or a transport you could then have:

```
headers_add = \  
  X-Spam-Scanned: ${primary_hostname} ${message_exim_id} \  
  ${hmac{md5}{SPAMSCAN_SECRET}\  
  ${primary_hostname},${message_exim_id},${h_message-id:}}
```

Then given a message, you can check where it was scanned by looking at the *X-Spam-Scanned:* header line. If you know the secret, you can check that this header line is authentic by recomputing the authentication code from the host name, message ID and the *Message-id:* header line. This can be done using Exim's **-be** option, or by other means, for example by using the *hmac\_md5\_hex()* function in Perl.

**`${if <condition> {<string1>}{<string2>}}`**

If *<condition>* is true, *<string1>* is expanded and replaces the whole item; otherwise *<string2>* is used. The available conditions are described in section 11.7 below. For example:

```
${if eq {$local_part}{postmaster} {yes}{no} }
```

The second string need not be present; if it is not and the condition is not true, the item is replaced with nothing. Alternatively, the word “fail” may be present instead of the second string (without any curly brackets). In this case, the expansion is forced to fail if the condition is not true (see section 11.4).

If both strings are omitted, the result is the string `true` if the condition is true, and the empty string if the condition is false. This makes it less cumbersome to write custom ACL and router conditions. For example, instead of

```
condition = ${if >{$acl_m4}{3}{true}{false}}
```

you can use

```
condition = ${if >{$acl_m4}{3}}
```

**`${imapfolder{<foldername>}}`**

This item converts a (possibly multilevel, or with non-ASCII characters) folder specification to a Maildir name for filesystem use. For information on internationalisation support see 59.2.

**`${length{<string1>}{<string2>}}`**

The **length** item is used to extract the initial portion of a string. Both strings are expanded, and the first one must yield a number, *<n>*, say. If you are using a fixed value for the number, that is, if *<string1>* does not change when expanded, you can use the simpler operator notation that avoids some of the braces:

```
${length_<n>:<string>}
```

The result of this item is either the first *<n>* characters or the whole of *<string2>*, whichever is the shorter. Do not confuse **length** with **strlen**, which gives the length of a string.

**`${listextract{<number>}{<string1>}{<string2>}{<string3>}}`**

The *<number>* argument must consist entirely of decimal digits, apart from an optional leading minus, and leading and trailing white space (which is ignored).

After expansion, *<string1>* is interpreted as a list, colon-separated by default, but the separator can be changed in the usual way.

The first field of the list is numbered one. If the number is negative, the fields are counted from the end of the list, with the rightmost one numbered -1. The numbered element of the list is extracted and placed in *\$value*, then *<string2>* is expanded as the result.

If the modulus of the number is zero or greater than the number of fields in the string, the result is the expansion of *<string3>*.

For example:

```
${listextract{2}{x:42:99}}
```

yields “42”, and

```
${listextract{-3}{<, x, 42, 99, & Mailer, , /bin/bash}{result: $value}}
```

yields “result: 42”.

If {<string3>} is omitted, an empty string is used for string3. If {<string2>} is also omitted, the value that was extracted is used. You can use `fail` instead of {<string3>} as in a string extract.

**`\${lookup}{<key>} <search type> {<file>} {<string1>} {<string2>}}`**

This is the first of one of two different types of lookup item, which are both described in the next item.

**`\${lookup} <search type> {<query>} {<string1>} {<string2>}}`**

The two forms of lookup item specify data lookups in files and databases, as discussed in chapter 9. The first form is used for single-key lookups, and the second is used for query-style lookups. The <key>, <file>, and <query> strings are expanded before use.

If there is any white space in a lookup item which is part of a filter command, a retry or rewrite rule, a routing rule for the *manualroute* router, or any other place where white space is significant, the lookup item must be enclosed in double quotes. The use of data lookups in users’ filter files may be locked out by the system administrator.

If the lookup succeeds, <string1> is expanded and replaces the entire item. During its expansion, the variable *\$value* contains the data returned by the lookup. Afterwards it reverts to the value it had previously (at the outer level it is empty). If the lookup fails, <string2> is expanded and replaces the entire item. If {<string2>} is omitted, the replacement is the empty string on failure. If <string2> is provided, it can itself be a nested lookup, thus providing a mechanism for looking up a default value when the original lookup fails.

If a nested lookup is used as part of <string1>, *\$value* contains the data for the outer lookup while the parameters of the second lookup are expanded, and also while <string2> of the second lookup is expanded, should the second lookup fail. Instead of {<string2>} the word “fail” can appear, and in this case, if the lookup fails, the entire expansion is forced to fail (see section 11.4). If both {<string1>} and {<string2>} are omitted, the result is the looked up value in the case of a successful lookup, and nothing in the case of failure.

For single-key lookups, the string “partial” is permitted to precede the search type in order to do partial matching, and \* or \*@ may follow a search type to request default lookups if the key does not match (see sections 9.6 and 9.7 for details).

If a partial search is used, the variables *\$1* and *\$2* contain the wild and non-wild parts of the key during the expansion of the replacement text. They return to their previous values at the end of the lookup item.

This example looks up the postmaster alias in the conventional alias file:

```
${lookup {postmaster} lsearch {/etc/aliases} {$value}}
```

This example uses NIS+ to look up the full name of the user corresponding to the local part of an address, forcing the expansion to fail if it is not found:

```
${lookup nisplus {[name=$local_part],passwd.org_dir:gcoss} \
{$value}fail}
```

**`\${map}{<string1>}{<string2>}}`**

After expansion, <string1> is interpreted as a list, colon-separated by default, but the separator can be changed in the usual way. For each item in this list, its value is placed in *\$item*, and then <string2> is expanded and added to the output as an item in a new list. The separator used for the output list is the same as the one used for the input, but a separator setting is not included in the output. For example:

```
${map{a:b:c}{[$item]}} ${map{<- x-y-z}{($item)}}
```

expands to `[a] : [b] : [c] (x) - (y) - (z)`. At the end of the expansion, the value of `$item` is restored to what it was before. See also the **filter** and **reduce** expansion items.

**`$(nhash{<string1>}{<string2>}{<string3>})`**

The three strings are expanded; the first two must yield numbers. Call them `<n>` and `<m>`. If you are using fixed values for these numbers, that is, if `<string1>` and `<string2>` do not change when they are expanded, you can use the simpler operator notation that avoids some of the braces:

```
$(nhash_<n>_<m>:<string>)
```

The second number is optional (in both notations). If there is only one number, the result is a number in the range `0-<n>-1`. Otherwise, the string is processed by a div/mod hash function that returns two numbers, separated by a slash, in the ranges `0 to <n>-1` and `0 to <m>-1`, respectively. For example,

```
$(nhash{8}{64}{supercalifragilisticexpialidocious})
```

returns the string “6/33”.

**`$(perl{<subroutine>}{<arg>}{<arg>}...}`**

This item is available only if Exim has been built to include an embedded Perl interpreter. The subroutine name and the arguments are first separately expanded, and then the Perl subroutine is called with those arguments. No additional arguments need be given; the maximum number permitted, including the name of the subroutine, is nine.

The return value of the subroutine is inserted into the expanded string, unless the return value is **undef**. In that case, the expansion fails in the same way as an explicit “fail” on a lookup item. The return value is a scalar. Whatever you return is evaluated in a scalar context. For example, if you return the name of a Perl vector, the return value is the size of the vector, not its contents.

If the subroutine exits by calling Perl’s **die** function, the expansion fails with the error message that was passed to **die**. More details of the embedded Perl facility are given in chapter 12.

The *redirect* router has an option called **forbid\_filter\_perl** which locks out the use of this expansion item in filter files.

**`$(prvs{<address>}{<secret>}{<keynumber>})`**

The first argument is a complete email address and the second is secret keystring. The third argument, specifying a key number, is optional. If absent, it defaults to 0. The result of the expansion is a prvs-signed email address, to be typically used with the **return\_path** option on an *smtp* transport as part of a bounce address tag validation (BATV) scheme. For more discussion and an example, see section 43.51.

**`$(prvscheck{<address>}{<secret>}{<string>})`**

This expansion item is the complement of the **prvs** item. It is used for checking prvs-signed addresses. If the expansion of the first argument does not yield a syntactically valid prvs-signed address, the whole item expands to the empty string. When the first argument does expand to a syntactically valid prvs-signed address, the second argument is expanded, with the prvs-decoded version of the address and the key number extracted from the address in the variables `$prvscheck_address` and `$prvscheck_keynum`, respectively.

These two variables can be used in the expansion of the second argument to retrieve the secret. The validity of the prvs-signed address is then checked against the secret. The result is stored in the variable `$prvscheck_result`, which is empty for failure or “1” for success.

The third argument is optional; if it is missing, it defaults to an empty string. This argument is now expanded. If the result is an empty string, the result of the expansion is the decoded version of the address. This is the case whether or not the signature was valid. Otherwise, the result of the expansion is the expansion of the third argument.

All three variables can be used in the expansion of the third argument. However, once the expansion is complete, only `$prvscheck_result` remains set. For more discussion and an example, see section 43.51.

**`${readfile}<file name>{<eol string>}`**

The file name and end-of-line string are first expanded separately. The file is then read, and its contents replace the entire item. All newline characters in the file are replaced by the end-of-line string if it is present. Otherwise, newlines are left in the string. String expansion is not applied to the contents of the file. If you want this, you must wrap the item in an **expand** operator. If the file cannot be read, the string expansion fails.

The *redirect* router has an option called **forbid\_filter\_readfile** which locks out the use of this expansion item in filter files.

**`${readsocket}<name>{<request>}{<options>}{<eol string>}{<fail string>}`**

This item inserts data from a Unix domain or TCP socket into the expanded string. The minimal way of using it uses just two arguments, as in these examples:

```
${readsocket{/socket/name}{request string}}
${readsocket{inet:some.host:1234}{request string}}
```

For a Unix domain socket, the first substring must be the path to the socket. For an Internet socket, the first substring must contain `inet:` followed by a host name or IP address, followed by a colon and a port, which can be a number or the name of a TCP port in */etc/services*. An IP address may optionally be enclosed in square brackets. This is best for IPv6 addresses. For example:

```
${readsocket{inet:[::1]:1234}{request string}}
```

Only a single host name may be given, but if looking it up yields more than one IP address, they are each tried in turn until a connection is made. For both kinds of socket, Exim makes a connection, writes the request string unless it is an empty string; and no terminating NUL is ever sent) and reads from the socket until an end-of-file is read. A timeout of 5 seconds is applied. Additional, optional arguments extend what can be done. Firstly, you can vary the timeout. For example:

```
${readsocket{/socket/name}{request string}{3s}}
```

The third argument is a list of options, of which the first element is the timeout and must be present if the argument is given. Further elements are options of form *name=value*. One option type is currently recognised, defining whether (the default) or not a shutdown is done on the connection after sending the request. Example, to not do so (preferred, eg. by some webservers):

```
${readsocket{/socket/name}{request string}{3s:shutdown=no}}
```

A fourth argument allows you to change any newlines that are in the data that is read, in the same way as for **readfile** (see above). This example turns them into spaces:

```
${readsocket{inet:127.0.0.1:3294}{request string}{3s}{ }}
```

As with all expansions, the substrings are expanded before the processing happens. Errors in these sub-expansions cause the expansion to fail. In addition, the following errors can occur:

- Failure to create a socket file descriptor;
- Failure to connect the socket;
- Failure to write the request string;
- Timeout on reading from the socket.

By default, any of these errors causes the expansion to fail. However, if you supply a fifth substring, it is expanded and used when any of the above errors occurs. For example:

```
${readsocket{/socket/name}{request string}{3s}{\n}\
{socket failure}}
```

You can test for the existence of a Unix domain socket by wrapping this expansion in `${if exists}`, but there is a race condition between that test and the actual opening of the socket, so it is safer to use the fifth argument if you want to be absolutely sure of avoiding an expansion error for a non-existent Unix domain socket, or a failure to connect to an Internet socket.

The *redirect* router has an option called **forbid\_filter\_readsocket** which locks out the use of this expansion item in filter files.

**`${reduce{<string1>}{<string2>}{<string3>}}`**

This operation reduces a list to a single, scalar string. After expansion, *<string1>* is interpreted as a list, colon-separated by default, but the separator can be changed in the usual way. Then *<string2>* is expanded and assigned to the *\$value* variable. After this, each item in the *<string1>* list is assigned to *\$item* in turn, and *<string3>* is expanded for each of them. The result of that expansion is assigned to *\$value* before the next iteration. When the end of the list is reached, the final value of *\$value* is added to the expansion output. The **reduce** expansion item can be used in a number of ways. For example, to add up a list of numbers:

```
${reduce {<, 1,2,3>}{0}{{eval:$value+$item}}}
```

The result of that expansion would be 6. The maximum of a list of numbers can be found:

```
${reduce {3:0:9:4:6}{0}{{if >{$item}{$value}}{$item}{$value}}}}
```

At the end of a **reduce** expansion, the values of *\$item* and *\$value* are restored to what they were before. See also the **filter** and **map** expansion items.

**`$rheader_<header name>:`** or **`$rh_<header name>:`**

This item inserts “raw” header lines. It is described with the **header** expansion item above.

**`${run{<command> <args>}{<string1>}{<string2>}}`**

The command and its arguments are first expanded as one string. The string is split apart into individual arguments by spaces, and then the command is run in a separate process, but under the same uid and gid. As in other command executions from Exim, a shell is not used by default. If the command requires a shell, you must explicitly code it.

Since the arguments are split by spaces, when there is a variable expansion which has an empty result, it will cause the situation that the argument will simply be omitted when the program is actually executed by Exim. If the script/program requires a specific number of arguments and the expanded variable could possibly result in this empty expansion, the variable must be quoted. This is more difficult if the expanded variable itself could result in a string containing quotes, because it would interfere with the quotes around the command arguments. A possible guard against this is to wrap the variable in the **sg** operator to change any quote marks to some other character.

The standard input for the command exists, but is empty. The standard output and standard error are set to the same file descriptor. If the command succeeds (gives a zero return code) *<string1>* is expanded and replaces the entire item; during this expansion, the standard output/error from the command is in the variable *\$value*. If the command fails, *<string2>*, if present, is expanded and used. Once again, during the expansion, the standard output/error from the command is in the variable *\$value*.

If *<string2>* is absent, the result is empty. Alternatively, *<string2>* can be the word “fail” (not in braces) to force expansion failure if the command does not succeed. If both strings are omitted, the result is contents of the standard output/error on success, and nothing on failure.

The standard output/error of the command is put in the variable *\$value*. In this ACL example, the output of a command is logged for the admin to troubleshoot:

```
warn    condition      = ${run{/usr/bin/id}{yes}{no}}
        log_message    = Output of id: $value
```

If the command requires shell idioms, such as the **>** redirect operator, the shell must be invoked directly, such as with:

```
${run{/bin/bash -c "/usr/bin/id >/tmp/id"}{yes}{yes}}
```

The return code from the command is put in the variable *\$runrc*, and this remains set afterwards, so in a filter file you can do things like this:

```
if "${run{x y z}{}}$runrc" is 1 then ...
elif $runrc is 2 then ...
```

```
...
endif
```

If execution of the command fails (for example, the command does not exist), the return code is 127 – the same code that shells use for non-existent commands.

**Warning:** In a router or transport, you cannot assume the order in which option values are expanded, except for those preconditions whose order of testing is documented. Therefore, you cannot reliably expect to set *\$runrc* by the expansion of one option, and use it in another.

The *redirect* router has an option called **forbid\_filter\_run** which locks out the use of this expansion item in filter files.

**`$(sg{<subject>}{<regex>}{<replacement>})`**

This item works like Perl's substitution operator (s) with the global (/g) option; hence its name. However, unlike the Perl equivalent, Exim does not modify the subject string; instead it returns the modified string for insertion into the overall expansion. The item takes three arguments: the subject string, a regular expression, and a substitution string. For example:

```
$(sg{abcdefabcdef}{abc}{xyz})
```

yields "xyzdefxyzdef". Because all three arguments are expanded before use, if any \$, } or \ characters are required in the regular expression or in the substitution string, they have to be escaped. For example:

```
$(sg{abcdef}{^(...) (...) \${} {\$2\$1})
```

yields "defabc", and

```
$(sg{1=A 4=D 3=C}{\N(\d+)=\N}{K\$1=})
```

yields "K1=A K4=D K3=C". Note the use of \N to protect the contents of the regular expression from string expansion.

**`$(sort{<string>}{<comparator>}{<extractor>})`**

After expansion, <string> is interpreted as a list, colon-separated by default, but the separator can be changed in the usual way. The <comparator> argument is interpreted as the operator of a two-argument expansion condition. The numeric operators plus ge, gt, le, lt (and ~i variants) are supported. The comparison should return true when applied to two values if the first value should sort before the second value. The <extractor> expansion is applied repeatedly to elements of the list, the element being placed in *\$item*, to give values for comparison.

The item result is a sorted list, with the original list separator, of the list elements (in full) of the original.

Examples:

```
$(sort{3:2:1:4}{<}{ $item})
```

sorts a list of numbers, and

```
$(sort { ${lookup dnsdb{>:, ,mx=example.com}} } {< } { ${listextract{1}}{<,$item}} )
```

will sort an MX lookup into priority order.

**`$(substr{<string1>}{<string2>}{<string3>})`**

The three strings are expanded; the first two must yield numbers. Call them <n> and <m>. If you are using fixed values for these numbers, that is, if <string1> and <string2> do not change when they are expanded, you can use the simpler operator notation that avoids some of the braces:

```
$(substr_<n>_<m>:<string>)
```

The second number is optional (in both notations). If it is absent in the simpler format, the preceding underscore must also be omitted.

The **substr** item can be used to extract more general substrings than **length**. The first number, <n>, is a starting offset, and <m> is the length required. For example

```
$(substr{3}{2}{ $local_part})
```

If the starting offset is greater than the string length the result is the null string; if the length plus starting offset is greater than the string length, the result is the right-hand part of the string, starting from the given offset. The first character in the string has offset zero.

The **substr** expansion item can take negative offset values to count from the right-hand end of its operand. The last character is offset -1, the second-last is offset -2, and so on. Thus, for example,

```
${substr{-5}{2}{1234567}}
```

yields “34”. If the absolute value of a negative offset is greater than the length of the string, the substring starts at the beginning of the string, and the length is reduced by the amount of overshoot. Thus, for example,

```
${substr{-5}{2}{12}}
```

yields an empty string, but

```
${substr{-3}{2}{12}}
```

yields “1”.

When the second number is omitted from **substr**, the remainder of the string is taken if the offset is positive. If it is negative, all characters in the string preceding the offset point are taken. For example, an offset of -1 and no length, as in these semantically identical examples:

```
${substr_-1:abcde}
${substr{-1}{abcde}}
```

yields all but the last character of the string, that is, “abcd”.

**`${tr{<subject>}{<characters>}{<replacements>}}`**

This item does single-character translation on its subject string. The second argument is a list of characters to be translated in the subject string. Each matching character is replaced by the corresponding character from the replacement list. For example

```
${tr{abcdea}{ac}{13}}
```

yields 1b3de1. If there are duplicates in the second character string, the last occurrence is used. If the third string is shorter than the second, its last character is replicated. However, if it is empty, no translation takes place.

## 11.6 Expansion operators

For expansion items that perform transformations on a single argument string, the “operator” notation is used because it is simpler and uses fewer braces. The substring is first expanded before the operation is applied to it. The following operations can be performed:

**`${address:<string>}`**

The string is interpreted as an RFC 2822 address, as it might appear in a header line, and the effective address is extracted from it. If the string does not parse successfully, the result is empty.

**`${addresses:<string>}`**

The string (after expansion) is interpreted as a list of addresses in RFC 2822 format, such as can be found in a *To:* or *Cc:* header line. The operative address (*local-part@domain*) is extracted from each item, and the result of the expansion is a colon-separated list, with appropriate doubling of colons should any happen to be present in the email addresses. Syntactically invalid RFC2822 address items are omitted from the output.

It is possible to specify a character other than colon for the output separator by starting the string with > followed by the new separator character. For example:

```
${addresses:>& Chief <ceo@up.stairs>, sec@base.ment (dogsbody)}
```

expands to ceo@up.stairs&sec@base.ment. The string is expanded first, so if the expanded string starts with >, it may change the output separator unintentionally. This can be avoided by setting the output separator explicitly:



```
${addresses:>:$h_from:}
```

Compare the **address** (singular) expansion item, which extracts the working address from a single RFC2822 address. See the **filter**, **map**, and **reduce** items for ways of processing lists.

To clarify "list of addresses in RFC 2822 format" mentioned above, Exim follows a strict interpretation of header line formatting. Exim parses the bare, unquoted portion of an email address and if it finds a comma, treats it as an email address separator. For the example header line:

```
From: =?iso-8859-2?Q?Last=2C_First?= <user@example.com>
```

The first example below demonstrates that Q-encoded email addresses are parsed properly if it is given the raw header (in this example, `$rheader_from:`). It does not see the comma because it's still encoded as `"=2C"`. The second example below is passed the contents of `$header_from:`, meaning it gets de-mimed. Exim sees the decoded `"`, so it treats it as **two** email addresses. The third example shows that the presence of a comma is skipped when it is quoted.

```
# exim -be '${addresses:From: \
=?iso-8859-2?Q?Last=2C_First?= <user@example.com>}'
user@example.com
# exim -be '${addresses:From: Last, First <user@example.com>}'
Last:user@example.com
# exim -be '${addresses:From: "Last, First" <user@example.com>}'
user@example.com
```

### **\${base32:<digits>}**

The string must consist entirely of decimal digits. The number is converted to base 32 and output as a (empty, for zero) string of characters. Only lowercase letters are used.

### **\${base32d:<base-32 digits>}**

The string must consist entirely of base-32 digits. The number is converted to decimal and output as a string.

### **\${base62:<digits>}**

The string must consist entirely of decimal digits. The number is converted to base 62 and output as a string of six characters, including leading zeros. In the few operating environments where Exim uses base 36 instead of base 62 for its message identifiers (because those systems do not have case-sensitive file names), base 36 is used by this operator, despite its name. **Note:** Just to be absolutely clear: this is *not* base64 encoding.

### **\${base62d:<base-62 digits>}**

The string must consist entirely of base-62 digits, or, in operating environments where Exim uses base 36 instead of base 62 for its message identifiers, base-36 digits. The number is converted to decimal and output as a string.

### **\${base64:<string>}**

This operator converts a string into one that is base64 encoded.

If the string is a single variable of type certificate, returns the base64 encoding of the DER form of the certificate.

### **\${base64d:<string>}**

This operator converts a base64-encoded string into the un-coded form.

### **\${domain:<string>}**

The string is interpreted as an RFC 2822 address and the domain is extracted from it. If the string does not parse successfully, the result is empty.

### **\${escape:<string>}**

If the string contains any non-printing characters, they are converted to escape sequences starting with a backslash. Whether characters with the most significant bit set (so-called "8-bit characters") count as printing or not is controlled by the **print\_topbitchars** option.

**`${escape8bit:<string>}`**

If the string contains and characters with the most significant bit set, they are converted to escape sequences starting with a backslash. Backslashes and DEL characters are also converted.

**`${eval:<string>}`** and **`${eval10:<string>}`**

These items supports simple arithmetic and bitwise logical operations in expansion strings. The string (after expansion) must be a conventional arithmetic expression, but it is limited to basic arithmetic operators, bitwise logical operators, and parentheses. All operations are carried out using integer arithmetic. The operator priorities are as follows (the same as in the C programming language):

<i>highest:</i>	not (~), negate (-)
	multiply (*), divide (/), remainder (%)
	plus (+), minus (-)
	shift-left (<<), shift-right (>>)
	and (&)
	xor (^)
<i>lowest:</i>	or ( )

Binary operators with the same priority are evaluated from left to right. White space is permitted before or after operators.

For **eval**, numbers may be decimal, octal (starting with “0”) or hexadecimal (starting with “0x”). For **eval10**, all numbers are taken as decimal, even if they start with a leading zero; hexadecimal numbers are not permitted. This can be useful when processing numbers extracted from dates or times, which often do have leading zeros.

A number may be followed by “K”, “M” or “G” to multiply it by 1024, 1024\*1024 or 1024\*1024\*1024, respectively. Negative numbers are supported. The result of the computation is a decimal representation of the answer (without “K”, “M” or “G”). For example:

<code>\${eval:1+1}</code>	yields 2
<code>\${eval:1+2*3}</code>	yields 7
<code>\${eval:(1+2)*3}</code>	yields 9
<code>\${eval:2+42%5}</code>	yields 4
<code>\${eval:0xc&amp;5}</code>	yields 4
<code>\${eval:0xc 5}</code>	yields 13
<code>\${eval:0xc^5}</code>	yields 9
<code>\${eval:0xc&gt;&gt;1}</code>	yields 6
<code>\${eval:0xc&lt;&lt;1}</code>	yields 24
<code>\${eval:~255&amp;0x1234}</code>	yields 4608
<code>\${eval:- (~255&amp;0x1234)}</code>	yields -4608

As a more realistic example, in an ACL you might have

```
deny    message = Too many bad recipients
        condition = \
        ${if and { \
            {>{$rcpt_count}{10}} \
            { \
                < \
                {$recipients_count} \
                ${eval:$rcpt_count/2}} \
            } \
        }{yes}{no}}
```

The condition is true if there have been more than 10 RCPT commands and fewer than half of them have resulted in a valid recipient.

**`${expand:<string>}`**

The **expand** operator causes a string to be expanded for a second time. For example,

```
${expand:${lookup{$domain}dbm{/some/file}{$value}}}
```

first looks up a string in a file while expanding the operand for **expand**, and then re-expands what it has found.

**`\${from\_utf8:<string>}`**

The world is slowly moving towards Unicode, although there are no standards for email yet. However, other applications (including some databases) are starting to store data in Unicode, using UTF-8 encoding. This operator converts from a UTF-8 string to an ISO-8859-1 string. UTF-8 code values greater than 255 are converted to underscores. The input must be a valid UTF-8 string. If it is not, the result is an undefined sequence of bytes.

Unicode code points with values less than 256 are compatible with ASCII and ISO-8859-1 (also known as Latin-1). For example, character 169 is the copyright symbol in both cases, though the way it is encoded is different. In UTF-8, more than one byte is needed for characters with code values greater than 127, whereas ISO-8859-1 is a single-byte encoding (but thereby limited to 256 characters). This makes translation from UTF-8 to ISO-8859-1 straightforward.

**`\${hash\_<n>\_<m>:<string>}`**

The **hash** operator is a simpler interface to the hashing function that can be used when the two parameters are fixed numbers (as opposed to strings that change when expanded). The effect is the same as

```
`${hash{<n>}{<m>}}{<string>}`
```

See the description of the general **hash** item above for details. The abbreviation **h** can be used when **hash** is used as an operator.

**`\${hex2b64:<hexstring>}`**

This operator converts a hex string into one that is base64 encoded. This can be useful for processing the output of the MD5 and SHA-1 hashing functions.

**`\${hexquote:<string>}`**

This operator converts non-printable characters in a string into a hex escape form. Byte values between 33 (!) and 126 (~) inclusive are left as is, and other byte values are converted to `\xNN`, for example a byte value 127 is converted to `\x7f`.

**`\${ipv6denorm:<string>}`**

This expands an IPv6 address to a full eight-element colon-separated set of hex digits including leading zeroes. A trailing ipv4-style dotted-decimal set is converted to hex. Pure IPv4 addresses are converted to IPv4-mapped IPv6.

**`\${ipv6norm:<string>}`**

This converts an IPv6 address to canonical form. Leading zeroes of groups are omitted, and the longest set of zero-valued groups is replaced with a double colon. A trailing ipv4-style dotted-decimal set is converted to hex. Pure IPv4 addresses are converted to IPv4-mapped IPv6.

**`\${lc:<string>}`**

This forces the letters in the string into lower-case, for example:

```
`${lc:$local_part}`
```

**`\${length\_<number>:<string>}`**

The **length** operator is a simpler interface to the **length** function that can be used when the parameter is a fixed number (as opposed to a string that changes when expanded). The effect is the same as

```
`${length{<number>}}{<string>}`
```

See the description of the general **length** item above for details. Note that **length** is not the same as **strlen**. The abbreviation **l** can be used when **length** is used as an operator.

**`\${listcount:<string>}`**

The string is interpreted as a list and the number of items is returned.

**`\${listnamed:<name>}`** and **`\${listnamed\_<type>:<name>}`**

The name is interpreted as a named list and the content of the list is returned, expanding any referenced lists, re-quoting as needed for colon-separation. If the optional type is given it must be

one of "a", "d", "h" or "l" and selects address-, domain-, host- or localpart- lists to search among respectively. Otherwise all types are searched in an undefined order and the first matching list is returned.

**`\${local\_part}<string>`**

The string is interpreted as an RFC 2822 address and the local part is extracted from it. If the string does not parse successfully, the result is empty.

**`\${mask}<IP address>/<bit count>`**

If the form of the string to be operated on is not an IP address followed by a slash and an integer (that is, a network address in CIDR notation), the expansion fails. Otherwise, this operator converts the IP address to binary, masks off the least significant bits according to the bit count, and converts the result back to text, with mask appended. For example,

```
`${mask}:10.111.131.206/28`
```

returns the string "10.111.131.192/28". Since this operation is expected to be mostly used for looking up masked addresses in files, the result for an IPv6 address uses dots to separate components instead of colons, because colon terminates a key string in lsearch files. So, for example,

```
`${mask}:3ffe:ffff:836f:0a00:000a:0800:200a:c031/99`
```

returns the string

```
3ffe.ffff.836f.0a00.000a.0800.2000.0000/99
```

Letters in IPv6 addresses are always output in lower case.

**`\${md5}<string>`**

The **md5** operator computes the MD5 hash value of the string, and returns it as a 32-digit hexadecimal number, in which any letters are in lower case.

If the string is a single variable of type certificate, returns the MD5 hash fingerprint of the certificate.

**`\${nhash}<n>\_<m>:<string>`**

The **nhash** operator is a simpler interface to the numeric hashing function that can be used when the two parameters are fixed numbers (as opposed to strings that change when expanded). The effect is the same as

```
`${nhash}<n>_<m>:<string>`
```

See the description of the general **nhash** item above for details.

**`\${quote}<string>`**

The **quote** operator puts its argument into double quotes if it is an empty string or contains anything other than letters, digits, underscores, dots, and hyphens. Any occurrences of double quotes and backslashes are escaped with a backslash. Newlines and carriage returns are converted to \n and \r, respectively. For example,

```
`${quote}:ab"*"cd`
```

becomes

```
"ab\"*"\"cd"
```

The place where this is useful is when the argument is a substitution from a variable or a message header.

**`\${quote\_local\_part}<string>`**

This operator is like **quote**, except that it quotes the string only if required to do so by the rules of RFC 2822 for quoting local parts. For example, a plus sign would not cause quoting (but it would for **quote**). If you are creating a new email address from the contents of *\$local\_part* (or any other unknown data), you should always use this operator.

**`\${quote\_<lookup-type>}<string>`**

This operator applies lookup-specific quoting rules to the string. Each query-style lookup type has its own quoting rules which are described with the lookups in chapter 9. For example,

```
${quote_ldap:two * two}
```

returns

```
two%20%5C2A%20two
```

For single-key lookup types, no quoting is ever necessary and this operator yields an unchanged string.

### **\${randint:<n>}**

This operator returns a somewhat random number which is less than the supplied number and is at least 0. The quality of this randomness depends on how Exim was built; the values are not suitable for keying material. If Exim is linked against OpenSSL then `RAND_pseudo_bytes()` is used. If Exim is linked against GnuTLS then `gnutls_rnd(GNUTLS_RND_NONCE)` is used, for versions of GnuTLS with that function. Otherwise, the implementation may be `arc4random()`, `random()` seeded by `srandomdev()` or `srandom()`, or a custom implementation even weaker than `random()`.

### **\${reverse\_ip:<ipaddr>}**

This operator reverses an IP address; for IPv4 addresses, the result is in dotted-quad decimal form, while for IPv6 addresses the result is in dotted-nibble hexadecimal form. In both cases, this is the "natural" form for DNS. For example,

```
${reverse_ip:192.0.2.4}
${reverse_ip:2001:0db8:c42:9:1:abcd:192.0.2.127}
```

returns

```
4.2.0.192
f.7.2.0.0.0.0.c.d.c.b.a.1.0.0.0.9.0.0.0.2.4.c.0.8.b.d.0.1.0.0.2
```

### **\${rfc2047:<string>}**

This operator encodes text according to the rules of RFC 2047. This is an encoding that is used in header lines to encode non-ASCII characters. It is assumed that the input string is in the encoding specified by the **headers\_charset** option, which gets its default at build time. If the string contains only characters in the range 33–126, and no instances of the characters

```
? = ( ) < > @ , ; : \ " . [ ] _
```

it is not modified. Otherwise, the result is the RFC 2047 encoding of the string, using as many “encoded words” as necessary to encode all the characters.

### **\${rfc2047d:<string>}**

This operator decodes strings that are encoded as per RFC 2047. Binary zero bytes are replaced by question marks. Characters are converted into the character set defined by **headers\_charset**. Overlong RFC 2047 “words” are not recognized unless **check\_rfc2047\_length** is set false.

**Note:** If you use **\$header\_xxx:** (or **\$h\_xxx:**) to access a header line, RFC 2047 decoding is done automatically. You do not need to use this operator as well.

### **\${rxquote:<string>}**

The **rxquote** operator inserts a backslash before any non-alphanumeric characters in its argument. This is useful when substituting the values of variables or headers inside regular expressions.

### **\${sha1:<string>}**

The **sha1** operator computes the SHA-1 hash value of the string, and returns it as a 40-digit hexadecimal number, in which any letters are in upper case.

If the string is a single variable of type certificate, returns the SHA-1 hash fingerprint of the certificate.

### **\${sha256:<string>}**

The **sha256** operator computes the SHA-256 hash value of the string and returns it as a 64-digit hexadecimal number, in which any letters are in upper case.

If the string is a single variable of type certificate, returns the SHA-256 hash fingerprint of the certificate.

**`${sha3:<string>}`**

**`${sha3_<n>:<string>}`**

The **sha3** operator computes the SHA3-256 hash value of the string and returns it as a 64-digit hexadecimal number, in which any letters are in upper case.

If a number is appended, separated by an underbar, it specifies the output length. Values of 224, 256, 384 and 512 are accepted; with 256 being the default.

The **sha3** expansion item is only supported if Exim has been compiled with GnuTLS 3.5.0 or later.

**`${stat:<string>}`**

The string, after expansion, must be a file path. A call to the *stat()* function is made for this path. If *stat()* fails, an error occurs and the expansion fails. If it succeeds, the data from the stat replaces the item, as a series of *<name>=<value>* pairs, where the values are all numerical, except for the value of “smode”. The names are: “mode” (giving the mode as a 4-digit octal number), “smode” (giving the mode in symbolic format as a 10-character string, as for the *ls* command), “inode”, “device”, “links”, “uid”, “gid”, “size”, “atime”, “mtime”, and “ctime”. You can extract individual fields using the **extract** expansion item.

The use of the **stat** expansion in users’ filter files can be locked out by the system administrator.

**Warning:** The file size may be incorrect on 32-bit systems for files larger than 2GB.

**`${str2b64:<string>}`**

Now deprecated, a synonym for the **base64** expansion operator.

**`${strlen:<string>}`**

The item is replaced by the length of the expanded string, expressed as a decimal number. **Note:** Do not confuse **strlen** with **length**.

**`${substr_<start>_<length>:<string>}`**

The **substr** operator is a simpler interface to the **substr** function that can be used when the two parameters are fixed numbers (as opposed to strings that change when expanded). The effect is the same as

```
${substr{<start>}{<length>}{<string>}}
```

See the description of the general **substr** item above for details. The abbreviation **s** can be used when **substr** is used as an operator.

**`${time_eval:<string>}`**

This item converts an Exim time interval such as 2d4h5m into a number of seconds.

**`${time_interval:<string>}`**

The argument (after sub-expansion) must be a sequence of decimal digits that represents an interval of time as a number of seconds. It is converted into a number of larger units and output in Exim’s normal time format, for example, 1w3d4h2m6s.

**`${uc:<string>}`**

This forces the letters in the string into upper-case.

**`${utf8clean:<string>}`**

This replaces any invalid utf-8 sequence in the string by the character ?.

**`${utf8_domain_to_alabel:<string>}`**

**`${utf8_domain_from_alabel:<string>}`**

**`${utf8_localpart_to_alabel:<string>}`**

**`${utf8_localpart_from_alabel:<string>}`**

These convert EAI mail name components between UTF-8 and a-label forms. For information on internationalisation support see 59.1.

## 11.7 Expansion conditions

The following conditions are available for testing by the **`${if}`** construct while expanding strings:

**!<condition>**

Preceding any condition with an exclamation mark negates the result of the condition.

**<symbolic operator> {<string1>}{<string2>}**

There are a number of symbolic operators for doing numeric comparisons. They are:

=	equal
==	equal
>	greater
>=	greater or equal
<	less
<=	less or equal

For example:

```
${if >{$message_size}{10M} ...}
```

Note that the general negation operator provides for inequality testing. The two strings must take the form of optionally signed decimal integers, optionally followed by one of the letters “K”, “M” or “G” (in either upper or lower case), signifying multiplication by 1024, 1024\*1024 or 1024\*1024\*1024, respectively. As a special case, the numerical value of an empty string is taken as zero.

In all cases, a relative comparator OP is testing if <string1> OP <string2>; the above example is checking if *\$message\_size* is larger than 10M, not if 10M is larger than *\$message\_size*.

**acl {{<name>}{<arg1>}{<arg2>}...}**

The name and zero to nine argument strings are first expanded separately. The expanded arguments are assigned to the variables *\$acl\_arg1* to *\$acl\_arg9* in order. Any unused are made empty. The variable *\$acl\_narg* is set to the number of arguments. The named ACL (see chapter 43) is called and may use the variables; if another acl expansion is used the values are restored after it returns. If the ACL sets a value using a “message =” modifier the variable *\$value* becomes the result of the expansion, otherwise it is empty. If the ACL returns accept the condition is true; if deny, false. If the ACL returns defer the result is a forced-fail.

**bool {<string>}**

This condition turns a string holding a true or false representation into a boolean state. It parses “true”, “false”, “yes” and “no” (case-insensitively); also integer numbers map to true if non-zero, false if zero. An empty string is treated as false. Leading and trailing whitespace is ignored; thus a string consisting only of whitespace is false. All other string values will result in expansion failure.

When combined with ACL variables, this expansion condition will let you make decisions in one place and act on those decisions in another place. For example:

```
${if bool{$acl_m_privileged_sender} ...}
```

**bool\_lax {<string>}**

Like **bool**, this condition turns a string into a boolean state. But where **bool** accepts a strict set of strings, **bool\_lax** uses the same loose definition that the Router **condition** option uses. The empty string and the values “false”, “no” and “0” map to false, all others map to true. Leading and trailing whitespace is ignored.

Note that where “bool{00}” is false, “bool\_lax{00}” is true.

**crypteq {<string1>}{<string2>}**

This condition is included in the Exim binary if it is built to support any authentication mechanisms (see chapter 33). Otherwise, it is necessary to define **SUPPORT\_CRYPTEQ** in *Local/Makefile* to get **crypteq** included in the binary.

The **crypteq** condition has two arguments. The first is encrypted and compared against the second, which is already encrypted. The second string may be in the LDAP form for storing encrypted strings, which starts with the encryption type in curly brackets, followed by the data. If the second string does not begin with “{” it is assumed to be encrypted with *crypt()* or *crypt16()* (see below), since such strings cannot begin with “{”. Typically this will be a field from a password file. An example of an encrypted string in LDAP form is:

```
{md5}CY9rzUYh03PK3k6DJie09g==
```

If such a string appears directly in an expansion, the curly brackets have to be quoted, because they are part of the expansion syntax. For example:

```
${if crypteq {test}{\{md5\}CY9rzUYh03PK3k6DJie09g==}{yes}{no}}
```

The following encryption types (whose names are matched case-independently) are supported:

- **{md5}** computes the MD5 digest of the first string, and expresses this as printable characters to compare with the remainder of the second string. If the length of the comparison string is 24, Exim assumes that it is base64 encoded (as in the above example). If the length is 32, Exim assumes that it is a hexadecimal encoding of the MD5 digest. If the length not 24 or 32, the comparison fails.
- **{sha1}** computes the SHA-1 digest of the first string, and expresses this as printable characters to compare with the remainder of the second string. If the length of the comparison string is 28, Exim assumes that it is base64 encoded. If the length is 40, Exim assumes that it is a hexadecimal encoding of the SHA-1 digest. If the length is not 28 or 40, the comparison fails.
- **{crypt}** calls the *crypt()* function, which traditionally used to use only the first eight characters of the password. However, in modern operating systems this is no longer true, and in many cases the entire password is used, whatever its length.
- **{crypt16}** calls the *crypt16()* function, which was originally created to use up to 16 characters of the password in some operating systems. Again, in modern operating systems, more characters may be used.

Exim has its own version of *crypt16()*, which is just a double call to *crypt()*. For operating systems that have their own version, setting `HAVE_CRYPT16` in *Local/Makefile* when building Exim causes it to use the operating system version instead of its own. This option is set by default in the OS-dependent *Makefile* for those operating systems that are known to support *crypt16()*.

Some years after Exim's *crypt16()* was implemented, a user discovered that it was not using the same algorithm as some operating systems' versions. It turns out that as well as *crypt16()* there is a function called *bigcrypt()* in some operating systems. This may or may not use the same algorithm, and both of them may be different to Exim's built-in *crypt16()*.

However, since there is now a move away from the traditional *crypt()* functions towards using SHA1 and other algorithms, tidying up this area of Exim is seen as very low priority.

If you do not put an encryption type (in curly brackets) in a **crypteq** comparison, the default is usually either **{crypt}** or **{crypt16}**, as determined by the setting of `DEFAULT_CRYPT` in *Local/Makefile*. The default default is **{crypt}**. Whatever the default, you can always use either function by specifying it explicitly in curly brackets.

#### **def:**<variable name>

The **def** condition must be followed by the name of one of the expansion variables defined in section 11.9. The condition is true if the variable does not contain the empty string. For example:

```
${if def:sender_ident {from $sender_ident}}
```

Note that the variable name is given without a leading **\$** character. If the variable does not exist, the expansion fails.

#### **def:header\_<header name>:** or **def:h\_<header name>:**

This condition is true if a message is being processed and the named header exists in the message. For example,

```
${if def:header_reply-to:{$h_reply-to:}{$h_from:}}
```

**Note:** No **\$** appears before **header\_** or **h\_** in the condition, and the header name must be terminated by a colon if white space does not follow.



**eq** {<string1>}{<string2>}

**eqi** {<string1>}{<string2>}

The two substrings are first expanded. The condition is true if the two resulting strings are identical. For **eq** the comparison includes the case of letters, whereas for **eqi** the comparison is case-independent.

**exists** {<file name>}

The substring is first expanded and then interpreted as an absolute path. The condition is true if the named file (or directory) exists. The existence test is done by calling the *stat()* function. The use of the **exists** test in users' filter files may be locked out by the system administrator.

**first\_delivery**

This condition, which has no data, is true during a message's first delivery attempt. It is false during any subsequent delivery attempts.

**forall**{<a list>}{<a condition>}

**forany**{<a list>}{<a condition>}

These conditions iterate over a list. The first argument is expanded to form the list. By default, the list separator is a colon, but it can be changed by the normal method. The second argument is interpreted as a condition that is to be applied to each item in the list in turn. During the interpretation of the condition, the current list item is placed in a variable called *\$item*.

- For **forany**, interpretation stops if the condition is true for any item, and the result of the whole condition is true. If the condition is false for all items in the list, the overall condition is false.
- For **forall**, interpretation stops if the condition is false for any item, and the result of the whole condition is false. If the condition is true for all items in the list, the overall condition is true.

Note that negation of **forany** means that the condition must be false for all items for the overall condition to succeed, and negation of **forall** means that the condition must be false for at least one item. In this example, the list separator is changed to a comma:

```
$ {if forany{<, $recipients}{match{$item}{^user3@}}{yes}{no}}
```

The value of *\$item* is saved and restored while **forany** or **forall** is being processed, to enable these expansion items to be nested.

To scan a named list, expand it with the **listnamed** operator.

**ge** {<string1>}{<string2>}

**gei** {<string1>}{<string2>}

The two substrings are first expanded. The condition is true if the first string is lexically greater than or equal to the second string. For **ge** the comparison includes the case of letters, whereas for **gei** the comparison is case-independent.

**gt** {<string1>}{<string2>}

**gti** {<string1>}{<string2>}

The two substrings are first expanded. The condition is true if the first string is lexically greater than the second string. For **gt** the comparison includes the case of letters, whereas for **gti** the comparison is case-independent.

**inlist** {<string1>}{<string2>}

**inlisti** {<string1>}{<string2>}

Both strings are expanded; the second string is treated as a list of simple strings; if the first string is a member of the second, then the condition is true.

These are simpler to use versions of the more powerful **forany** condition. Examples, and the **forany** equivalents:

```
$ {if inlist{needle}{foo:needle:bar}}
$ {if forany{foo:needle:bar}{eq{$item}{needle}}}
$ {if inlisti{Needle}{fOo:NeeDLE:bAr}}
$ {if forany{fOo:NeeDLE:bAr}{eqi{$item}{Needle}}}
```

**isip** {<string>}  
**isip4** {<string>}  
**isip6** {<string>}

The substring is first expanded, and then tested to see if it has the form of an IP address. Both IPv4 and IPv6 addresses are valid for **isip**, whereas **isip4** and **isip6** test specifically for IPv4 or IPv6 addresses.

For an IPv4 address, the test is for four dot-separated components, each of which consists of from one to three digits. For an IPv6 address, up to eight colon-separated components are permitted, each containing from one to four hexadecimal digits. There may be fewer than eight components if an empty component (adjacent colons) is present. Only one empty component is permitted.

**Note:** The checks used to be just on the form of the address; actual numerical values were not considered. Thus, for example, 999.999.999.999 passed the IPv4 check. This is no longer the case.

The main use of these tests is to distinguish between IP addresses and host names, or between IPv4 and IPv6 addresses. For example, you could use

```
${if isip4{$sender_host_address}}...
```

to test which IP version an incoming SMTP connection is using.

**ldapauth** {<ldap query>}

This condition supports user authentication using LDAP. See section 9.14 for details of how to use LDAP in lookups and the syntax of queries. For this use, the query must contain a user name and password. The query itself is not used, and can be empty. The condition is true if the password is not empty, and the user name and password are accepted by the LDAP server. An empty password is rejected without calling LDAP because LDAP binds with an empty password are considered anonymous regardless of the username, and will succeed in most configurations. See chapter 33 for details of SMTP authentication, and chapter 34 for an example of how this can be used.

**le** {<string1>}{<string2>}

**lei** {<string1>}{<string2>}

The two substrings are first expanded. The condition is true if the first string is lexically less than or equal to the second string. For **le** the comparison includes the case of letters, whereas for **lei** the comparison is case-independent.

**lt** {<string1>}{<string2>}

**lti** {<string1>}{<string2>}

The two substrings are first expanded. The condition is true if the first string is lexically less than the second string. For **lt** the comparison includes the case of letters, whereas for **lti** the comparison is case-independent.

**match** {<string1>}{<string2>}

The two substrings are first expanded. The second is then treated as a regular expression and applied to the first. Because of the pre-expansion, if the regular expression contains dollar, or backslash characters, they must be escaped. Care must also be taken if the regular expression contains braces (curly brackets). A closing brace must be escaped so that it is not taken as a premature termination of <string2>. The easiest approach is to use the \N feature to disable expansion of the regular expression. For example,

```
${if match {$local_part}{\N^d{3}\N}}...
```

If the whole expansion string is in double quotes, further escaping of backslashes is also required.

The condition is true if the regular expression match succeeds. The regular expression is not required to begin with a circumflex metacharacter, but if there is no circumflex, the expression is not anchored, and it may match anywhere in the subject, not just at the start. If you want the pattern to match at the end of the subject, you must include the \$ metacharacter at an appropriate point.

At the start of an **if** expansion the values of the numeric variable substitutions \$1 etc. are remembered. Obeying a **match** condition that succeeds causes them to be reset to the substrings of that condition and they will have these values during the expansion of the success string. At the end of

the **if** expansion, the previous values are restored. After testing a combination of conditions using **or**, the subsequent values of the numeric variables are those of the condition that succeeded.

**match\_address** {<string1>}{<string2>}

See **match\_local\_part**.

**match\_domain** {<string1>}{<string2>}

See **match\_local\_part**.

**match\_ip** {<string1>}{<string2>}

This condition matches an IP address to a list of IP address patterns. It must be followed by two argument strings. The first (after expansion) must be an IP address or an empty string. The second (not expanded) is a restricted host list that can match only an IP address, not a host name. For example:

```
{if match_ip{$sender_host_address}{1.2.3.4:5.6.7.8}{...}{...}}
```

The specific types of host list item that are permitted in the list are:

- An IP address, optionally with a CIDR mask.
- A single asterisk, which matches any IP address.
- An empty item, which matches only if the IP address is empty. This could be useful for testing for a locally submitted message or one from specific hosts in a single test such as

```
{if match_ip{$sender_host_address}{:4.3.2.1:...}{...}{...}}
```

where the first item in the list is the empty string.

- The item @[] matches any of the local host's interface addresses.
- Single-key lookups are assumed to be like "net-" style lookups in host lists, even if net- is not specified. There is never any attempt to turn the IP address into a host name. The most common type of linear search for **match\_ip** is likely to be **iplsearch**, in which the file can contain CIDR masks. For example:

```
{if match_ip{$sender_host_address}{iplsearch;/some/file}...
```

It is of course possible to use other kinds of lookup, and in such a case, you do need to specify the net- prefix if you want to specify a specific address mask, for example:

```
{if match_ip{$sender_host_address}{net24-dbm;/some/file}...
```

However, unless you are combining a **match\_ip** condition with others, it is just as easy to use the fact that a lookup is itself a condition, and write:

```
{lookup{$mask:$sender_host_address/24}dbm{/a/file}...
```

Note that <string2> is not itself subject to string expansion, unless Exim was built with the EXPAND\_LISTMATCH\_RHS option.

Consult section 10.11 for further details of these patterns.

**match\_local\_part** {<string1>}{<string2>}

This condition, together with **match\_address** and **match\_domain**, make it possible to test domain, address, and local part lists within expansions. Each condition requires two arguments: an item and a list to match. A trivial example is:

```
{if match_domain{a.b.c}{x.y.z:a.b.c:p.q.r}{yes}{no}}
```

In each case, the second argument may contain any of the allowable items for a list of the appropriate type. Also, because the second argument is a standard form of list, it is possible to refer to a named list. Thus, you can use conditions like this:

```
{if match_domain{$domain}{+local_domains}{...}}
```

For address lists, the matching starts off caselessly, but the +caseful item can be used, as in all address lists, to cause subsequent items to have their local parts matched casefully. Domains are always matched caselessly.

Note that `<string2>` is not itself subject to string expansion, unless Exim was built with the `EXPAND_LISTMATCH_RHS` option.

**Note:** Host lists are *not* supported in this way. This is because hosts have two identities: a name and an IP address, and it is not clear how to specify cleanly how such a test would work. However, IP addresses can be matched using **match\_ip**.

**pam** {<string1>:<string2>:...}

*Pluggable Authentication Modules* (<http://www.kernel.org/pub/linux/libs/pam/>) are a facility that is available in the latest releases of Solaris and in some GNU/Linux distributions. The Exim support, which is intended for use in conjunction with the SMTP AUTH command, is available only if Exim is compiled with

```
SUPPORT_PAM=yes
```

in *Local/Makefile*. You probably need to add **-lpam** to `EXTRALIBS`, and in some releases of GNU/Linux **-ldl** is also needed.

The argument string is first expanded, and the result must be a colon-separated list of strings. Leading and trailing white space is ignored. The PAM module is initialized with the service name “exim” and the user name taken from the first item in the colon-separated data string (<string1>). The remaining items in the data string are passed over in response to requests from the authentication function. In the simple case there will only be one request, for a password, so the data consists of just two strings.

There can be problems if any of the strings are permitted to contain colon characters. In the usual way, these have to be doubled to avoid being taken as separators. If the data is being inserted from a variable, the **sg** expansion item can be used to double any existing colons. For example, the configuration of a LOGIN authenticator might contain this setting:

```
server_condition = ${if pam{$auth1:${sg{$auth2}}{:}{:}}{}}
```

For a PLAIN authenticator you could use:

```
server_condition = ${if pam{$auth2:${sg{$auth3}}{:}{:}}{}}
```

In some operating systems, PAM authentication can be done only from a process running as root. Since Exim is running as the Exim user when receiving messages, this means that PAM cannot be used directly in those systems. A patched version of the *pam\_unix* module that comes with the Linux PAM package is available from [http://www.e-admin.de/pam\\_exim/](http://www.e-admin.de/pam_exim/). The patched module allows one special uid/gid combination, in addition to root, to authenticate. If you build the patched module to allow the Exim user and group, PAM can then be used from an Exim authenticator.

**pwcheck** {<string1>:<string2>}

This condition supports user authentication using the Cyrus *pwcheck* daemon. This is one way of making it possible for passwords to be checked by a process that is not running as root. **Note:** The use of *pwcheck* is now deprecated. Its replacement is *saslauthd* (see below).

The *pwcheck* support is not included in Exim by default. You need to specify the location of the *pwcheck* daemon’s socket in *Local/Makefile* before building Exim. For example:

```
CYRUS_PWCHECK_SOCKET=/var/pwcheck/pwcheck
```

You do not need to install the full Cyrus software suite in order to use the *pwcheck* daemon. You can compile and install just the daemon alone from the Cyrus SASL library. Ensure that *exim* is the only user that has access to the */var/pwcheck* directory.

The **pwcheck** condition takes one argument, which must be the user name and password, separated by a colon. For example, in a LOGIN authenticator configuration, you might have this:

```
server_condition = ${if pwcheck{$auth1:$auth2}}
```

Again, for a PLAIN authenticator configuration, this would be:

```
server_condition = ${if pwcheck{$auth2:$auth3}}
```

### **queue\_running**

This condition, which has no data, is true during delivery attempts that are initiated by queue runner processes, and false otherwise.

### **radius** {<authentication string>}

Radius authentication (RFC 2865) is supported in a similar way to PAM. You must set `RADIUS_CONFIG_FILE` in *Local/Makefile* to specify the location of the Radius client configuration file in order to build Exim with Radius support.

With just that one setting, Exim expects to be linked with the **radiusclient** library, using the original API. If you are using release 0.4.0 or later of this library, you need to set

```
RADIUS_LIB_TYPE=RADIUSCLIENTNEW
```

in *Local/Makefile* when building Exim. You can also link Exim with the **libradius** library that comes with FreeBSD. To do this, set

```
RADIUS_LIB_TYPE=RADLIB
```

in *Local/Makefile*, in addition to setting `RADIUS_CONFIGURE_FILE`. You may also have to supply a suitable setting in `EXTRALIBS` so that the Radius library can be found when Exim is linked.

The string specified by `RADIUS_CONFIG_FILE` is expanded and passed to the Radius client library, which calls the Radius server. The condition is true if the authentication is successful. For example:

```
server_condition = ${if radius{<arguments>}}
```

### **saslauthd** {{<user>}{<password>}{<service>}{<realm>}}

This condition supports user authentication using the Cyrus *saslauthd* daemon. This replaces the older *pwcheck* daemon, which is now deprecated. Using this daemon is one way of making it possible for passwords to be checked by a process that is not running as root.

The *saslauthd* support is not included in Exim by default. You need to specify the location of the *saslauthd* daemon's socket in *Local/Makefile* before building Exim. For example:

```
CYRUS_SASLAUTHD_SOCKET=/var/state/saslauthd/mux
```

You do not need to install the full Cyrus software suite in order to use the *saslauthd* daemon. You can compile and install just the daemon alone from the Cyrus SASL library.

Up to four arguments can be supplied to the **saslauthd** condition, but only two are mandatory. For example:

```
server_condition = ${if saslauthd{{${auth1}}${auth2}}}
```

The service and the realm are optional (which is why the arguments are enclosed in their own set of braces). For details of the meaning of the service and realm, and how to run the daemon, consult the Cyrus documentation.

## **11.8 Combining expansion conditions**

Several conditions can be tested at once by combining them using the **and** and **or** combination conditions. Note that **and** and **or** are complete conditions on their own, and precede their lists of sub-conditions. Each sub-condition must be enclosed in braces within the overall braces that contain the list. No repetition of **if** is used.

### **or** {{<cond1>}{<cond2>}...}

The sub-conditions are evaluated from left to right. The condition is true if any one of the sub-conditions is true. For example,

```
${if or {{eq{$local_part}{spqr}}{eq{$domain}{testing.com}}}}...
```

When a true sub-condition is found, the following ones are parsed but not evaluated. If there are several “match” sub-conditions the values of the numeric variables afterwards are taken from the first one that succeeds.

**and** {{<cond1>}}{<cond2>}...

The sub-conditions are evaluated from left to right. The condition is true if all of the sub-conditions are true. If there are several “match” sub-conditions, the values of the numeric variables afterwards are taken from the last one. When a false sub-condition is found, the following ones are parsed but not evaluated.

## 11.9 Expansion variables

This section contains an alphabetical list of all the expansion variables. Some of them are available only when Exim is compiled with specific options such as support for TLS or the content scanning extension.

*\$0, \$1, etc*

When a **match** expansion condition succeeds, these variables contain the captured substrings identified by the regular expression during subsequent processing of the success string of the containing **if** expansion item. In the expansion condition case they do not retain their values afterwards; in fact, their previous values are restored at the end of processing an **if** item. The numerical variables may also be set externally by some other matching process which precedes the expansion of the string. For example, the commands available in Exim filter files include an **if** command with its own regular expression matching condition.

*\$acl\_arg1, \$acl\_arg2, etc*

Within an acl condition, expansion condition or expansion item any arguments are copied to these variables, any unused variables being made empty.

*\$acl\_c...*

Values can be placed in these variables by the **set** modifier in an ACL. They can be given any name that starts with *\$acl\_c* and is at least six characters long, but the sixth character must be either a digit or an underscore. For example: *\$acl\_c5*, *\$acl\_c\_mycount*. The values of the *\$acl\_c...* variables persist throughout the lifetime of an SMTP connection. They can be used to pass information between ACLs and between different invocations of the same ACL. When a message is received, the values of these variables are saved with the message, and can be accessed by filters, routers, and transports during subsequent delivery.

*\$acl\_m...*

These variables are like the *\$acl\_c...* variables, except that their values are reset after a message has been received. Thus, if several messages are received in one SMTP connection, *\$acl\_m...* values are not passed on from one message to the next, as *\$acl\_c...* values are. The *\$acl\_m...* variables are also reset by MAIL, RSET, EHLO, HELO, and after starting a TLS session. When a message is received, the values of these variables are saved with the message, and can be accessed by filters, routers, and transports during subsequent delivery.

*\$acl\_narg*

Within an acl condition, expansion condition or expansion item this variable has the number of arguments.

*\$acl\_verify\_message*

After an address verification has failed, this variable contains the failure message. It retains its value for use in subsequent modifiers. The message can be preserved by coding like this:

```
warn !verify = sender
    set acl_m0 = $acl_verify_message
```

You can use *\$acl\_verify\_message* during the expansion of the **message** or **log\_message** modifiers, to include information about the verification failure.

*\$address\_data*

This variable is set by means of the **address\_data** option in routers. The value then remains with the address while it is processed by subsequent routers and eventually a transport. If the transport is handling multiple addresses, the value from the first address is used. See chapter 15 for more details. **Note:** The contents of *\$address\_data* are visible in user filter files.

If *\$address\_data* is set when the routers are called from an ACL to verify a recipient address, the final value is still in the variable for subsequent conditions and modifiers of the ACL statement. If routing the address caused it to be redirected to just one address, the child address is also routed as part of the verification, and in this case the final value of *\$address\_data* is from the child's routing.

If *\$address\_data* is set when the routers are called from an ACL to verify a sender address, the final value is also preserved, but this time in *\$sender\_address\_data*, to distinguish it from data from a recipient address.

In both cases (recipient and sender verification), the value does not persist after the end of the current ACL statement. If you want to preserve these values for longer, you can save them in ACL variables.

#### *\$address\_file*

When, as a result of aliasing, forwarding, or filtering, a message is directed to a specific file, this variable holds the name of the file when the transport is running. At other times, the variable is empty. For example, using the default configuration, if user **r2d2** has a *.forward* file containing

```
/home/r2d2/savemail
```

then when the *address\_file* transport is running, *\$address\_file* contains the text string */home/r2d2/savemail*. For Sieve filters, the value may be "inbox" or a relative folder name. It is then up to the transport configuration to generate an appropriate absolute path to the relevant file.

#### *\$address\_pipe*

When, as a result of aliasing or forwarding, a message is directed to a pipe, this variable holds the pipe command when the transport is running.

#### *\$auth1* – *\$auth3*

These variables are used in SMTP authenticators (see chapters 34–41). Elsewhere, they are empty.

#### *\$authenticated\_id*

When a server successfully authenticates a client it may be configured to preserve some of the authentication information in the variable *\$authenticated\_id* (see chapter 33). For example, a user/password authenticator configuration might preserve the user name for use in the routers. Note that this is not the same information that is saved in *\$sender\_host\_authenticated*. When a message is submitted locally (that is, not over a TCP connection) the value of *\$authenticated\_id* is normally the login name of the calling process. However, a trusted user can override this by means of the **-oMai** command line option.

#### *\$authenticated\_fail\_id*

When an authentication attempt fails, the variable *\$authenticated\_fail\_id* will contain the failed authentication id. If more than one authentication id is attempted, it will contain only the last one. The variable is available for processing in the ACL's, generally the quit or notquit ACL. A message to a local recipient could still be accepted without requiring authentication, which means this variable could also be visible in all of the ACL's as well.

#### *\$authenticated\_sender*

When acting as a server, Exim takes note of the AUTH= parameter on an incoming SMTP MAIL command if it believes the sender is sufficiently trusted, as described in section 33.2. Unless the data is the string "<>", it is set as the authenticated sender of the message, and the value is available during delivery in the *\$authenticated\_sender* variable. If the sender is not trusted, Exim accepts the syntax of AUTH=, but ignores the data.

When a message is submitted locally (that is, not over a TCP connection), the value of *\$authenticated\_sender* is an address constructed from the login name of the calling process and *\$qualify\_domain*, except that a trusted user can override this by means of the **-oMas** command line option.

#### *\$authentication\_failed*

This variable is set to "1" in an Exim server if a client issues an AUTH command that does not succeed. Otherwise it is set to "0". This makes it possible to distinguish between "did not try to authenticate" (*\$sender\_host\_authenticated* is empty and *\$authentication\_failed* is set to "0") and

“tried to authenticate but failed” (*\$sender\_host\_authenticated* is empty and *\$authentication\_failed* is set to “1”). Failure includes any negative response to an AUTH command, including (for example) an attempt to use an undefined mechanism.

#### *\$av\_failed*

This variable is available when Exim is compiled with the content-scanning extension. It is set to “0” by default, but will be set to “1” if any problem occurs with the virus scanner (specified by **av\_scanner**) during the ACL malware condition.

#### *\$body\_linecount*

When a message is being received or delivered, this variable contains the number of lines in the message’s body. See also *\$message\_linecount*.

#### *\$body\_zerocount*

When a message is being received or delivered, this variable contains the number of binary zero bytes (ASCII NULs) in the message’s body.

#### *\$bounce\_recipient*

This is set to the recipient address of a bounce message while Exim is creating it. It is useful if a customized bounce message text file is in use (see chapter 49).

#### *\$bounce\_return\_size\_limit*

This contains the value set in the **bounce\_return\_size\_limit** option, rounded up to a multiple of 1000. It is useful when a customized error message text file is in use (see chapter 49).

#### *\$caller\_gid*

The real group id under which the process that called Exim was running. This is not the same as the group id of the originator of a message (see *\$originator\_gid*). If Exim re-execs itself, this variable in the new incarnation normally contains the Exim gid.

#### *\$caller\_uid*

The real user id under which the process that called Exim was running. This is not the same as the user id of the originator of a message (see *\$originator\_uid*). If Exim re-execs itself, this variable in the new incarnation normally contains the Exim uid.

#### *\$callout\_address*

After a callout for verification, spamd or malware daemon service, the address that was connected to.

#### *\$compile\_number*

The building process for Exim keeps a count of the number of times it has been compiled. This serves to distinguish different compilations of the same version of the program.

#### *\$config\_dir*

The directory name of the main configuration file. That is, the content of *\$config\_file* with the last component stripped. The value does not contain the trailing slash. If *\$config\_file* does not contain a slash, *\$config\_dir* is “.”.

#### *\$config\_file*

The name of the main configuration file Exim is using.

#### *\$dkim\_verify\_status*

Results

of

DKIM

verification.

For details see chapter 57.



*\$dkim\_cur\_signer*  
*\$dkim\_verify\_reason*  
*\$dkim\_domain*  
*\$dkim\_identity*  
*\$dkim\_selector*  
*\$dkim\_algo*  
*\$dkim\_canon\_body*  
*\$dkim\_canon\_headers*  
*\$dkim\_copiedheaders*  
*\$dkim\_bodysize*  
*\$dkim\_created*  
*\$dkim\_expires*  
*\$dkim\_headernames*  
*\$dkim\_key\_testing*  
*\$dkim\_key\_nosubdomains*  
*\$dkim\_key\_srvtype*  
*\$dkim\_key\_granularity*  
*\$dkim\_key\_notes*  
*\$dkim\_key\_length*

These variables are only available within the DKIM ACL. For details see chapter 57.

*\$dkim\_signers*

When a message has been received this variable contains a colon-separated list of signer domains and identities for the message. For details see chapter 57.

*\$dnslist\_domain*  
*\$dnslist\_matched*  
*\$dnslist\_text*  
*\$dnslist\_value*

When a DNS (black) list lookup succeeds, these variables are set to contain the following data from the lookup: the list's domain name, the key that was looked up, the contents of any associated TXT record, and the value from the main A record. See section 43.32 for more details.

*\$domain*

When an address is being routed, or delivered on its own, this variable contains the domain. Uppercase letters in the domain are converted into lower case for *\$domain*.

Global address rewriting happens when a message is received, so the value of *\$domain* during routing and delivery is the value after rewriting. *\$domain* is set during user filtering, but not during system filtering, because a message may have many recipients and the system filter is called just once.

When more than one address is being delivered at once (for example, several RCPT commands in one SMTP delivery), *\$domain* is set only if they all have the same domain. Transports can be restricted to handling only one domain at a time if the value of *\$domain* is required at transport time – this is the default for local transports. For further details of the environment in which local transports are run, see chapter 23.

At the end of a delivery, if all deferred addresses have the same domain, it is set in *\$domain* during the expansion of **delay\_warning\_condition**.

The *\$domain* variable is also used in some other circumstances:

- When an ACL is running for a RCPT command, *\$domain* contains the domain of the recipient address. The domain of the *sender* address is in *\$sender\_address\_domain* at both MAIL time and at RCPT time. *\$domain* is not normally set during the running of the MAIL ACL. However, if the sender address is verified with a callout during the MAIL ACL, the sender domain is placed in *\$domain* during the expansions of **hosts**, **interface**, and **port** in the *smtp* transport.
- When a rewrite item is being processed (see chapter 31), *\$domain* contains the domain portion of the address that is being rewritten; it can be used in the expansion of the replacement address, for example, to rewrite domains by file lookup.

- With one important exception, whenever a domain list is being scanned, *\$domain* contains the subject domain. **Exception:** When a domain list in a **sender\_domains** condition in an ACL is being processed, the subject domain is in *\$sender\_address\_domain* and not in *\$domain*. It works this way so that, in a RCPT ACL, the sender domain list can be dependent on the recipient domain (which is what is in *\$domain* at this time).
- When the **smtp\_etrn\_command** option is being expanded, *\$domain* contains the complete argument of the ETRN command (see section 48.8).

#### *\$domain\_data*

When the **domains** option on a router matches a domain by means of a lookup, the data read by the lookup is available during the running of the router as *\$domain\_data*. In addition, if the driver routes the address to a transport, the value is available in that transport. If the transport is handling multiple addresses, the value from the first address is used.

*\$domain\_data* is also set when the **domains** condition in an ACL matches a domain by means of a lookup. The data read by the lookup is available during the rest of the ACL statement. In all other situations, this variable expands to nothing.

#### *\$exim\_gid*

This variable contains the numerical value of the Exim group id.

#### *\$exim\_path*

This variable contains the path to the Exim binary.

#### *\$exim\_uid*

This variable contains the numerical value of the Exim user id.

#### *\$exim\_version*

This variable contains the version string of the Exim build. The first character is a major version number, currently 4. Then after a dot, the next group of digits is a minor version number. There may be other characters following the minor version.

#### *\$header\_<name>*

This is not strictly an expansion variable. It is expansion syntax for inserting the message header line with the given name. Note that the name must be terminated by colon or white space, because it may contain a wide variety of characters. Note also that braces must *not* be used.

#### *\$headers\_added*

Within an ACL this variable contains the headers added so far by the ACL modifier **add\_header** (section 43.24). The headers are a newline-separated list.

#### *\$home*

When the **check\_local\_user** option is set for a router, the user's home directory is placed in *\$home* when the check succeeds. In particular, this means it is set during the running of users' filter files. A router may also explicitly set a home directory for use by a transport; this can be overridden by a setting on the transport itself.

When running a filter test via the **-bf** option, *\$home* is set to the value of the environment variable HOME, which is subject to the **keep\_environment** and **add\_environment** main config options.

#### *\$host*

If a router assigns an address to a transport (any transport), and passes a list of hosts with the address, the value of *\$host* when the transport starts to run is the name of the first host on the list. Note that this applies both to local and remote transports.

For the *smtp* transport, if there is more than one host, the value of *\$host* changes as the transport works its way through the list. In particular, when the *smtp* transport is expanding its options for encryption using TLS, or for specifying a transport filter (see chapter 24), *\$host* contains the name of the host to which it is connected.

When used in the client part of an authenticator configuration (see chapter 33), *\$host* contains the name of the server to which the client is connected.

### *\$host\_address*

This variable is set to the remote host's IP address whenever *\$host* is set for a remote connection. It is also set to the IP address that is being checked when the **ignore\_target\_hosts** option is being processed.

### *\$host\_data*

If a **hosts** condition in an ACL is satisfied by means of a lookup, the result of the lookup is made available in the *\$host\_data* variable. This allows you, for example, to do things like this:

```
deny hosts = net-lsearch;/some/file
message = $host_data
```

### *\$host\_lookup\_deferred*

This variable normally contains "0", as does *\$host\_lookup\_failed*. When a message comes from a remote host and there is an attempt to look up the host's name from its IP address, and the attempt is not successful, one of these variables is set to "1".

- If the lookup receives a definite negative response (for example, a DNS lookup succeeded, but no records were found), *\$host\_lookup\_failed* is set to "1".
- If there is any kind of problem during the lookup, such that Exim cannot tell whether or not the host name is defined (for example, a timeout for a DNS lookup), *\$host\_lookup\_deferred* is set to "1".

Looking up a host's name from its IP address consists of more than just a single reverse lookup. Exim checks that a forward lookup of at least one of the names it receives from a reverse lookup yields the original IP address. If this is not the case, Exim does not accept the looked up name(s), and *\$host\_lookup\_failed* is set to "1". Thus, being able to find a name from an IP address (for example, the existence of a PTR record in the DNS) is not sufficient on its own for the success of a host name lookup. If the reverse lookup succeeds, but there is a lookup problem such as a timeout when checking the result, the name is not accepted, and *\$host\_lookup\_deferred* is set to "1". See also *\$sender\_host\_name*.

### *\$host\_lookup\_failed*

See *\$host\_lookup\_deferred*.

### *\$host\_port*

This variable is set to the remote host's TCP port whenever *\$host* is set for an outbound connection.

### *\$initial\_cwd*

This variable contains the full path name of the initial working directory of the current Exim process. This may differ from the current working directory, as Exim changes this to "/" during early startup, and to *\$spool\_directory* later.

### *\$inode*

The only time this variable is set is while expanding the **directory\_file** option in the *appendfile* transport. The variable contains the inode number of the temporary file which is about to be renamed. It can be used to construct a unique name for the file.

### *\$interface\_address*

This is an obsolete name for *\$received\_ip\_address*.

### *\$interface\_port*

This is an obsolete name for *\$received\_port*.

### *\$item*

This variable is used during the expansion of **forall** and **forany** conditions (see section 11.7), and **filter**, **map**, and **reduce** items (see section 11.7). In other circumstances, it is empty.

### *\$ldap\_dn*

This variable, which is available only when Exim is compiled with LDAP support, contains the DN from the last entry in the most recently successful LDAP lookup.

### *\$load\_average*

This variable contains the system load average, multiplied by 1000 so that it is an integer. For example, if the load average is 0.21, the value of the variable is 210. The value is recomputed every time the variable is referenced.

### *\$local\_part*

When an address is being routed, or delivered on its own, this variable contains the local part. When a number of addresses are being delivered together (for example, multiple RCPT commands in an SMTP session), *\$local\_part* is not set.

Global address rewriting happens when a message is received, so the value of *\$local\_part* during routing and delivery is the value after rewriting. *\$local\_part* is set during user filtering, but not during system filtering, because a message may have many recipients and the system filter is called just once.

If a local part prefix or suffix has been recognized, it is not included in the value of *\$local\_part* during routing and subsequent delivery. The values of any prefix or suffix are in *\$local\_part\_prefix* and *\$local\_part\_suffix*, respectively.

When a message is being delivered to a file, pipe, or autoreply transport as a result of aliasing or forwarding, *\$local\_part* is set to the local part of the parent address, not to the file name or command (see *\$address\_file* and *\$address\_pipe*).

When an ACL is running for a RCPT command, *\$local\_part* contains the local part of the recipient address.

When a rewrite item is being processed (see chapter 31), *\$local\_part* contains the local part of the address that is being rewritten; it can be used in the expansion of the replacement address, for example.

In all cases, all quoting is removed from the local part. For example, for both the addresses

```
"abc:xyz"@test.example
abc\:xyz@test.example
```

the value of *\$local\_part* is

```
abc:xyz
```

If you use *\$local\_part* to create another address, you should always wrap it inside a quoting operator. For example, in a *redirect* router you could have:

```
data = ${quote_local_part:$local_part}@new.domain.example
```

**Note:** The value of *\$local\_part* is normally lower cased. If you want to process local parts in a case-dependent manner in a router, you can set the **caseful\_local\_part** option (see chapter 15).

### *\$local\_part\_data*

When the **local\_parts** option on a router matches a local part by means of a lookup, the data read by the lookup is available during the running of the router as *\$local\_part\_data*. In addition, if the driver routes the address to a transport, the value is available in that transport. If the transport is handling multiple addresses, the value from the first address is used.

*\$local\_part\_data* is also set when the **local\_parts** condition in an ACL matches a local part by means of a lookup. The data read by the lookup is available during the rest of the ACL statement. In all other situations, this variable expands to nothing.

### *\$local\_part\_prefix*

When an address is being routed or delivered, and a specific prefix for the local part was recognized, it is available in this variable, having been removed from *\$local\_part*.

### *\$local\_part\_suffix*

When an address is being routed or delivered, and a specific suffix for the local part was recognized, it is available in this variable, having been removed from *\$local\_part*.

#### *\$local\_scan\_data*

This variable contains the text returned by the *local\_scan()* function when a message is received. See chapter 45 for more details.

#### *\$local\_user\_gid*

See *\$local\_user\_uid*.

#### *\$local\_user\_uid*

This variable and *\$local\_user\_gid* are set to the uid and gid after the **check\_local\_user** router precondition succeeds. This means that their values are available for the remaining preconditions (**senders**, **require\_files**, and **condition**), for the **address\_data** expansion, and for any router-specific expansions. At all other times, the values in these variables are (uid\_t) (-1) and (gid\_t) (-1), respectively.

#### *\$localhost\_number*

This contains the expanded value of the **localhost\_number** option. The expansion happens after the main options have been read.

#### *\$log\_inodes*

The number of free inodes in the disk partition where Exim's log files are being written. The value is recalculated whenever the variable is referenced. If the relevant file system does not have the concept of inodes, the value of is -1. See also the **check\_log\_inodes** option.

#### *\$log\_space*

The amount of free space (as a number of kilobytes) in the disk partition where Exim's log files are being written. The value is recalculated whenever the variable is referenced. If the operating system does not have the ability to find the amount of free space (only true for experimental systems), the space value is -1. See also the **check\_log\_space** option.

#### *\$lookup\_dnssec\_authenticated*

This variable is set after a DNS lookup done by a dnsdb lookup expansion, dnslookup router or smtp transport. It will be empty if *DNSSEC* was not requested, “no” if the result was not labelled as authenticated data and “yes” if it was. Results that are labelled as authoritative answer that match the **dns\_trust\_aa** configuration variable count also as authenticated data.

#### *\$mailstore\_basename*

This variable is set only when doing deliveries in “mailstore” format in the *appendfile* transport. During the expansion of the **mailstore\_prefix**, **mailstore\_suffix**, **message\_prefix**, and **message\_suffix** options, it contains the basename of the files that are being written, that is, the name without the “.tmp”, “.env”, or “.msg” suffix. At all other times, this variable is empty.

#### *\$malware\_name*

This variable is available when Exim is compiled with the content-scanning extension. It is set to the name of the virus that was found when the ACL **malware** condition is true (see section 44.1).

#### *\$max\_received\_linelength*

This variable contains the number of bytes in the longest line that was received as part of the message, not counting the line termination character(s).

It is not valid if the **spool\_files\_wireformat** option is used.

#### *\$message\_age*

This variable is set at the start of a delivery attempt to contain the number of seconds since the message was received. It does not change during a single delivery attempt.

#### *\$message\_body*

This variable contains the initial portion of a message's body while it is being delivered, and is intended mainly for use in filter files. The maximum number of characters of the body that are put into the variable is set by the **message\_body\_visible** configuration option; the default is 500.

By default, newlines are converted into spaces in *\$message\_body*, to make it easier to search for phrases that might be split over a line break. However, this can be disabled by setting **message\_body\_newlines** to be true. Binary zeros are always converted into spaces.

### *\$message\_body\_end*

This variable contains the final portion of a message's body while it is being delivered. The format and maximum size are as for *\$message\_body*.

### *\$message\_body\_size*

When a message is being delivered, this variable contains the size of the body in bytes. The count starts from the character after the blank line that separates the body from the header. Newlines are included in the count. See also *\$message\_size*, *\$body\_linecount*, and *\$body\_zerocount*.

If the spool file is wireformat (see the **spool\_files\_wireformat** main option) the CRLF line-terminators are included in the count.

### *\$message\_exim\_id*

When a message is being received or delivered, this variable contains the unique message id that is generated and used by Exim to identify the message. An id is not created for a message until after its header has been successfully received. **Note:** This is *not* the contents of the *Message-ID:* header line; it is the local id that Exim assigns to the message, for example: 1BXTIK-0001yO-VA.

### *\$message\_headers*

This variable contains a concatenation of all the header lines when a message is being processed, except for lines added by routers or transports. The header lines are separated by newline characters. Their contents are decoded in the same way as a header line that is inserted by **bheader**.

### *\$message\_headers\_raw*

This variable is like *\$message\_headers* except that no processing of the contents of header lines is done.

### *\$message\_id*

This is an old name for *\$message\_exim\_id*. It is now deprecated.

### *\$message\_linecount*

This variable contains the total number of lines in the header and body of the message. Compare *\$body\_linecount*, which is the count for the body only. During the DATA and content-scanning ACLs, *\$message\_linecount* contains the number of lines received. Before delivery happens (that is, before filters, routers, and transports run) the count is increased to include the *Received:* header line that Exim standardly adds, and also any other header lines that are added by ACLs. The blank line that separates the message header from the body is not counted.

As with the special case of *\$message\_size*, during the expansion of the appendfile transport's *maildir\_tag* option in maildir format, the value of *\$message\_linecount* is the precise size of the number of newlines in the file that has been written (minus one for the blank line between the header and the body).

Here is an example of the use of this variable in a DATA ACL:

```
deny message    = Too many lines in message header
condition = \
    ${if <{250}}{${eval:$message_linecount - $body_linecount}}}
```

In the MAIL and RCPT ACLs, the value is zero because at that stage the message has not yet been received.

This variable is not valid if the **spool\_files\_wireformat** option is used.

### *\$message\_size*

When a message is being processed, this variable contains its size in bytes. In most cases, the size includes those headers that were received with the message, but not those (such as *Envelope-to:*) that are added to individual deliveries as they are written. However, there is one special case: during the expansion of the **maildir\_tag** option in the *appendfile* transport while doing a delivery in maildir format, the value of *\$message\_size* is the precise size of the file that has been written. See also *\$message\_body\_size*, *\$body\_linecount*, and *\$body\_zerocount*.

While running a per message ACL (mail/rcpt/predata), *\$message\_size* contains the size supplied on the MAIL command, or -1 if no size was given. The value may not, of course, be truthful.

### *\$mime\_xxx*

A number of variables whose names start with *\$mime* are available when Exim is compiled with the content-scanning extension. For details, see section 44.4.

### *\$n0 – \$n9*

These variables are counters that can be incremented by means of the **add** command in filter files.

### *\$original\_domain*

When a top-level address is being processed for delivery, this contains the same value as *\$domain*. However, if a “child” address (for example, generated by an alias, forward, or filter file) is being processed, this variable contains the domain of the original address (lower cased). This differs from *\$parent\_domain* only when there is more than one level of aliasing or forwarding. When more than one address is being delivered in a single transport run, *\$original\_domain* is not set.

If a new address is created by means of a **deliver** command in a system filter, it is set up with an artificial “parent” address. This has the local part *system-filter* and the default qualify domain.

### *\$original\_local\_part*

When a top-level address is being processed for delivery, this contains the same value as *\$local\_part*, unless a prefix or suffix was removed from the local part, because *\$original\_local\_part* always contains the full local part. When a “child” address (for example, generated by an alias, forward, or filter file) is being processed, this variable contains the full local part of the original address.

If the router that did the redirection processed the local part case-insensitively, the value in *\$original\_local\_part* is in lower case. This variable differs from *\$parent\_local\_part* only when there is more than one level of aliasing or forwarding. When more than one address is being delivered in a single transport run, *\$original\_local\_part* is not set.

If a new address is created by means of a **deliver** command in a system filter, it is set up with an artificial “parent” address. This has the local part *system-filter* and the default qualify domain.

### *\$originator\_gid*

This variable contains the value of *\$caller\_gid* that was set when the message was received. For messages received via the command line, this is the gid of the sending user. For messages received by SMTP over TCP/IP, this is normally the gid of the Exim user.

### *\$originator\_uid*

The value of *\$caller\_uid* that was set when the message was received. For messages received via the command line, this is the uid of the sending user. For messages received by SMTP over TCP/IP, this is normally the uid of the Exim user.

### *\$parent\_domain*

This variable is similar to *\$original\_domain* (see above), except that it refers to the immediately preceding parent address.

### *\$parent\_local\_part*

This variable is similar to *\$original\_local\_part* (see above), except that it refers to the immediately preceding parent address.

### *\$pid*

This variable contains the current process id.

### *\$pipe\_addresses*

This is not an expansion variable, but is mentioned here because the string *\$pipe\_addresses* is handled specially in the command specification for the *pipe* transport (chapter 29) and in transport filters (described under **transport\_filter** in chapter 24). It cannot be used in general expansion strings, and provokes an “unknown variable” error if encountered.

### *\$primary\_hostname*

This variable contains the value set by **primary\_hostname** in the configuration file, or read by the *uname()* function. If *uname()* returns a single-component name, Exim calls *gethostbyname()* (or *getipnodebyname()* where available) in an attempt to acquire a fully qualified host name. See also *\$smtp\_active\_hostname*.

*\$proxy\_external\_address*

*\$proxy\_external\_port*

*\$proxy\_local\_address*

*\$proxy\_local\_port*

*\$proxy\_session*

These variables are only available when built with Proxy Protocol or SOCKS5 support. For details see chapter 58.1.

*\$prdr\_requested*

This variable is set to “yes” if PRDR was requested by the client for the current message, otherwise “no”.

*\$prvscheck\_address*

This variable is used in conjunction with the **prvscheck** expansion item, which is described in sections 11.5 and 43.51.

*\$prvscheck\_keynum*

This variable is used in conjunction with the **prvscheck** expansion item, which is described in sections 11.5 and 43.51.

*\$prvscheck\_result*

This variable is used in conjunction with the **prvscheck** expansion item, which is described in sections 11.5 and 43.51.

*\$qualify\_domain*

The value set for the **qualify\_domain** option in the configuration file.

*\$qualify\_recipient*

The value set for the **qualify\_recipient** option in the configuration file, or if not set, the value of *\$qualify\_domain*.

*\$queue\_name*

The name of the spool queue in use; empty for the default queue.

*\$rcpt\_count*

When a message is being received by SMTP, this variable contains the number of RCPT commands received for the current message. If this variable is used in a RCPT ACL, its value includes the current command.

*\$rcpt\_defer\_count*

When a message is being received by SMTP, this variable contains the number of RCPT commands in the current message that have previously been rejected with a temporary (4xx) response.

*\$rcpt\_fail\_count*

When a message is being received by SMTP, this variable contains the number of RCPT commands in the current message that have previously been rejected with a permanent (5xx) response.

*\$received\_count*

This variable contains the number of *Received:* header lines in the message, including the one added by Exim (so its value is always greater than zero). It is available in the DATA ACL, the non-SMTP ACL, and while routing and delivering.

*\$received\_for*

If there is only a single recipient address in an incoming message, this variable contains that address when the *Received:* header line is being built. The value is copied after recipient rewriting has happened, but before the *local\_scan()* function is run.

*\$received\_ip\_address*

As soon as an Exim server starts processing an incoming TCP/IP connection, this variable is set to the address of the local IP interface, and *\$received\_port* is set to the local port number. (The remote IP address and port are in *\$sender\_host\_address* and *\$sender\_host\_port*.) When testing with **-bh**, the port value is -1 unless it has been set using the **-oMi** command line option.

As well as being useful in ACLs (including the “connect” ACL), these variable could be used, for example, to make the file name for a TLS certificate depend on which interface and/or port is



being used for the incoming connection. The values of `$received_ip_address` and `$received_port` are saved with any messages that are received, thus making these variables available at delivery time. For outbound connections see `$sending_ip_address`.

#### `$received_port`

See `$received_ip_address`.

#### `$received_protocol`

When a message is being processed, this variable contains the name of the protocol by which it was received. Most of the names used by Exim are defined by RFCs 821, 2821, and 3848. They start with “smtp” (the client used HELO) or “esmt” (the client used EHLO). This can be followed by “s” for secure (encrypted) and/or “a” for authenticated. Thus, for example, if the protocol is set to “esmtpsa”, the message was received over an encrypted SMTP connection and the client was successfully authenticated.

Exim uses the protocol name “smtps” for the case when encryption is automatically set up on connection without the use of STARTTLS (see **tls\_on\_connect\_ports**), and the client uses HELO to initiate the encrypted SMTP session. The name “smtps” is also used for the rare situation where the client initially uses EHLO, sets up an encrypted connection using STARTTLS, and then uses HELO afterwards.

The **-oMr** option provides a way of specifying a custom protocol name for messages that are injected locally by trusted callers. This is commonly used to identify messages that are being re-injected after some kind of scanning.

#### `$received_time`

This variable contains the date and time when the current message was received, as a number of seconds since the start of the Unix epoch.

#### `$recipient_data`

This variable is set after an indexing lookup success in an ACL **recipients** condition. It contains the data from the lookup, and the value remains set until the next **recipients** test. Thus, you can do things like this:

```
require recipients = cdb*@;/some/file
deny      some further test involving $recipient_data
```

**Warning:** This variable is set only when a lookup is used as an indexing method in the address list, using the semicolon syntax as in the example above. The variable is not set for a lookup that is used as part of the string expansion that all such lists undergo before being interpreted.

#### `$recipient_verify_failure`

In an ACL, when a recipient verification fails, this variable contains information about the failure. It is set to one of the following words:

- “qualify”: The address was unqualified (no domain), and the message was neither local nor came from an exempted host.
- “route”: Routing failed.
- “mail”: Routing succeeded, and a callout was attempted; rejection occurred at or before the MAIL command (that is, on initial connection, HELO, or MAIL).
- “recipient”: The RCPT command in a callout was rejected.
- “postmaster”: The postmaster check in a callout was rejected.

The main use of this variable is expected to be to distinguish between rejections of MAIL and rejections of RCPT.

#### `$recipients`

This variable contains a list of envelope recipients for a message. A comma and a space separate the addresses in the replacement text. However, the variable is not generally available, to prevent exposure of Bcc recipients in unprivileged users’ filter files. You can use `$recipients` only in these cases:

- (1) In a system filter file.
- (2) In the ACLs associated with the DATA command and with non-SMTP messages, that is, the ACLs defined by **acl\_smtp\_predata**, **acl\_smtp\_data**, **acl\_smtp\_mime**, **acl\_not\_smtp\_start**, **acl\_not\_smtp**, and **acl\_not\_smtp\_mime**.
- (3) From within a *local\_scan()* function.

#### *\$recipients\_count*

When a message is being processed, this variable contains the number of envelope recipients that came with the message. Duplicates are not excluded from the count. While a message is being received over SMTP, the number increases for each accepted recipient. It can be referenced in an ACL.

#### *\$regex\_match\_string*

This variable is set to contain the matching regular expression after a **regex** ACL condition has matched (see section 44.5).

#### *\$regex1*, *\$regex2*, etc

When a **regex** or **mime\_regex** ACL condition succeeds, these variables contain the captured substrings identified by the regular expression.

#### *\$reply\_address*

When a message is being processed, this variable contains the contents of the *Reply-To:* header line if one exists and it is not empty, or otherwise the contents of the *From:* header line. Apart from the removal of leading white space, the value is not processed in any way. In particular, no RFC 2047 decoding or character code translation takes place.

#### *\$return\_path*

When a message is being delivered, this variable contains the return path – the sender field that will be sent as part of the envelope. It is not enclosed in <> characters. At the start of routing an address, *\$return\_path* has the same value as *\$sender\_address*, but if, for example, an incoming message to a mailing list has been expanded by a router which specifies a different address for bounce messages, *\$return\_path* subsequently contains the new bounce address, whereas *\$sender\_address* always contains the original sender address that was received with the message. In other words, *\$sender\_address* contains the incoming envelope sender, and *\$return\_path* contains the outgoing envelope sender.

#### *\$return\_size\_limit*

This is an obsolete name for *\$bounce\_return\_size\_limit*.

#### *\$router\_name*

During the running of a router this variable contains its name.

#### *\$runrc*

This variable contains the return code from a command that is run by the **\${run...}** expansion item.

**Warning:** In a router or transport, you cannot assume the order in which option values are expanded, except for those preconditions whose order of testing is documented. Therefore, you cannot reliably expect to set *\$runrc* by the expansion of one option, and use it in another.

#### *\$self\_hostname*

When an address is routed to a supposedly remote host that turns out to be the local host, what happens is controlled by the **self** generic router option. One of its values causes the address to be passed to another router. When this happens, *\$self\_hostname* is set to the name of the local host that the original router encountered. In other circumstances its contents are null.

#### *\$sender\_address*

When a message is being processed, this variable contains the sender's address that was received in the message's envelope. The case of letters in the address is retained, in both the local part and the domain. For bounce messages, the value of this variable is the empty string. See also *\$return\_path*.

### *\$sender\_address\_data*

If *\$address\_data* is set when the routers are called from an ACL to verify a sender address, the final value is preserved in *\$sender\_address\_data*, to distinguish it from data from a recipient address. The value does not persist after the end of the current ACL statement. If you want to preserve it for longer, you can save it in an ACL variable.

### *\$sender\_address\_domain*

The domain portion of *\$sender\_address*.

### *\$sender\_address\_local\_part*

The local part portion of *\$sender\_address*.

### *\$sender\_data*

This variable is set after a lookup success in an ACL **senders** condition or in a router **senders** option. It contains the data from the lookup, and the value remains set until the next **senders** test. Thus, you can do things like this:

```
require senders      = cdb*@;/some/file
deny      some further test involving $sender_data
```

**Warning:** This variable is set only when a lookup is used as an indexing method in the address list, using the semicolon syntax as in the example above. The variable is not set for a lookup that is used as part of the string expansion that all such lists undergo before being interpreted.

### *\$sender\_fullhost*

When a message is received from a remote host, this variable contains the host name and IP address in a single string. It ends with the IP address in square brackets, followed by a colon and a port number if the logging of ports is enabled. The format of the rest of the string depends on whether the host issued a HELO or EHLO SMTP command, and whether the host name was verified by looking up its IP address. (Looking up the IP address can be forced by the **host\_lookup** option, independent of verification.) A plain host name at the start of the string is a verified host name; if this is not present, verification either failed or was not requested. A host name in parentheses is the argument of a HELO or EHLO command. This is omitted if it is identical to the verified host name or to the host's IP address in square brackets.

### *\$sender\_helo\_dnssec*

This boolean variable is true if a successful HELO verification was done using DNS information the resolver library stated was authenticated data.

### *\$sender\_helo\_name*

When a message is received from a remote host that has issued a HELO or EHLO command, the argument of that command is placed in this variable. It is also set if HELO or EHLO is used when a message is received using SMTP locally via the **-bs** or **-bS** options.

### *\$sender\_host\_address*

When a message is received from a remote host using SMTP, this variable contains that host's IP address. For locally non-SMTP submitted messages, it is empty.

### *\$sender\_host\_authenticated*

This variable contains the name (not the public name) of the authenticator driver that successfully authenticated the client from which the message was received. It is empty if there was no successful authentication. See also *\$authenticated\_id*.

### *\$sender\_host\_dnssec*

If an attempt to populate *\$sender\_host\_name* has been made (by reference, **hosts\_lookup** or otherwise) then this boolean will have been set true if, and only if, the resolver library states that both the reverse and forward DNS were authenticated data. At all other times, this variable is false.

It is likely that you will need to coerce DNSSEC support on in the resolver library, by setting:

```
dns_dnssec_ok = 1
```

Exim does not perform DNSSEC validation itself, instead leaving that to a validating resolver (e.g. unbound, or bind with suitable configuration).

If you have changed **host\_lookup\_order** so that `bydns` is not the first mechanism in the list, then this variable will be false.

This requires that your system resolver library support EDNS0 (and that DNSSEC flags exist in the system headers). If the resolver silently drops all EDNS0 options, then this will have no effect. OpenBSD's `asr` resolver is known to currently ignore EDNS0, documented in CAVEATS of `asr_run(3)`.

#### *\$sender\_host\_name*

When a message is received from a remote host, this variable contains the host's name as obtained by looking up its IP address. For messages received by other means, this variable is empty.

If the host name has not previously been looked up, a reference to *\$sender\_host\_name* triggers a lookup (for messages from remote hosts). A looked up name is accepted only if it leads back to the original IP address via a forward lookup. If either the reverse or the forward lookup fails to find any data, or if the forward lookup does not yield the original IP address, *\$sender\_host\_name* remains empty, and *\$host\_lookup\_failed* is set to "1".

However, if either of the lookups cannot be completed (for example, there is a DNS timeout), *\$host\_lookup\_deferred* is set to "1", and *\$host\_lookup\_failed* remains set to "0".

Once *\$host\_lookup\_failed* is set to "1", Exim does not try to look up the host name again if there is a subsequent reference to *\$sender\_host\_name* in the same Exim process, but it does try again if *\$host\_lookup\_deferred* is set to "1".

Exim does not automatically look up every calling host's name. If you want maximum efficiency, you should arrange your configuration so that it avoids these lookups altogether. The lookup happens only if one or more of the following are true:

- A string containing *\$sender\_host\_name* is expanded.
- The calling host matches the list in **host\_lookup**. In the default configuration, this option is set to \*, so it must be changed if lookups are to be avoided. (In the code, the default for **host\_lookup** is unset.)
- Exim needs the host name in order to test an item in a host list. The items that require this are described in sections 10.13 and 10.17.
- The calling host matches **helo\_try\_verify\_hosts** or **helo\_verify\_hosts**. In this case, the host name is required to compare with the name quoted in any EHLO or HELO commands that the client issues.
- The remote host issues a EHLO or HELO command that quotes one of the domains in **helo\_lookup\_domains**. The default value of this option is

```
helo_lookup_domains = @ : @[]
```

which causes a lookup if a remote host (incorrectly) gives the server's name or IP address in an EHLO or HELO command.

#### *\$sender\_host\_port*

When a message is received from a remote host, this variable contains the port number that was used on the remote host.

#### *\$sender\_ident*

When a message is received from a remote host, this variable contains the identification received in response to an RFC 1413 request. When a message has been received locally, this variable contains the login name of the user that called Exim.

#### *\$sender\_rate\_xxx*

A number of variables whose names begin *\$sender\_rate\_* are set as part of the **ratelimit** ACL condition. Details are given in section 43.38.

#### *\$sender\_rcvhost*

This is provided specifically for use in *Received:* headers. It starts with either the verified host name (as obtained from a reverse DNS lookup) or, if there is no verified host name, the IP address

in square brackets. After that there may be text in parentheses. When the first item is a verified host name, the first thing in the parentheses is the IP address in square brackets, followed by a colon and a port number if port logging is enabled. When the first item is an IP address, the port is recorded as “port=xxxx” inside the parentheses.

There may also be items of the form “helo=xxxx” if HELO or EHLO was used and its argument was not identical to the real host name or IP address, and “ident=xxxx” if an RFC 1413 ident string is available. If all three items are present in the parentheses, a newline and tab are inserted into the string, to improve the formatting of the *Received:* header.

#### *\$sender\_verify\_failure*

In an ACL, when a sender verification fails, this variable contains information about the failure. The details are the same as for *\$recipient\_verify\_failure*.

#### *\$sending\_ip\_address*

This variable is set whenever an outgoing SMTP connection to another host has been set up. It contains the IP address of the local interface that is being used. This is useful if a host that has more than one IP address wants to take on different personalities depending on which one is being used. For incoming connections, see *\$received\_ip\_address*.

#### *\$sending\_port*

This variable is set whenever an outgoing SMTP connection to another host has been set up. It contains the local port that is being used. For incoming connections, see *\$received\_port*.

#### *\$smtp\_active\_hostname*

During an incoming SMTP session, this variable contains the value of the active host name, as specified by the **smtp\_active\_hostname** option. The value of *\$smtp\_active\_hostname* is saved with any message that is received, so its value can be consulted during routing and delivery.

#### *\$smtp\_command*

During the processing of an incoming SMTP command, this variable contains the entire command. This makes it possible to distinguish between HELO and EHLO in the HELO ACL, and also to distinguish between commands such as these:

```
MAIL FROM:<>
MAIL FROM: <>
```

For a MAIL command, extra parameters such as SIZE can be inspected. For a RCPT command, the address in *\$smtp\_command* is the original address before any rewriting, whereas the values in *\$local\_part* and *\$domain* are taken from the address after SMTP-time rewriting.

#### *\$smtp\_command\_argument*

While an ACL is running to check an SMTP command, this variable contains the argument, that is, the text that follows the command name, with leading white space removed. Following the introduction of *\$smtp\_command*, this variable is somewhat redundant, but is retained for backwards compatibility.

#### *\$smtp\_command\_history*

A comma-separated list (with no whitespace) of the most-recent SMTP commands received, in time-order left to right. Only a limited number of commands are remembered.

#### *\$smtp\_count\_at\_connection\_start*

This variable is set greater than zero only in processes spawned by the Exim daemon for handling incoming SMTP connections. The name is deliberately long, in order to emphasize what the contents are. When the daemon accepts a new connection, it increments this variable. A copy of the variable is passed to the child process that handles the connection, but its value is fixed, and never changes. It is only an approximation of how many incoming connections there actually are, because many other connections may come and go while a single connection is being processed. When a child process terminates, the daemon decrements its copy of the variable.

#### *\$sn0 – \$sn9*

These variables are copies of the values of the *\$n0 – \$n9* accumulators that were current at the end of the system filter file. This allows a system filter file to set values that can be tested in users’

filter files. For example, a system filter could set a value indicating how likely it is that a message is junk mail.

#### *\$spam\_xxx*

A number of variables whose names start with *\$spam* are available when Exim is compiled with the content-scanning extension. For details, see section 44.2.

#### *\$spool\_directory*

The name of Exim's spool directory.

#### *\$spool\_inodes*

The number of free inodes in the disk partition where Exim's spool files are being written. The value is recalculated whenever the variable is referenced. If the relevant file system does not have the concept of inodes, the value of is -1. See also the **check\_spool\_inodes** option.

#### *\$spool\_space*

The amount of free space (as a number of kilobytes) in the disk partition where Exim's spool files are being written. The value is recalculated whenever the variable is referenced. If the operating system does not have the ability to find the amount of free space (only true for experimental systems), the space value is -1. For example, to check in an ACL that there is at least 50 megabytes free on the spool, you could write:

```
condition = ${if > {$spool_space}{50000}}
```

See also the **check\_spool\_space** option.

#### *\$thisaddress*

This variable is set only during the processing of the **foranyaddress** command in a filter file. Its use is explained in the description of that command, which can be found in the separate document entitled *Exim's interfaces to mail filtering*.

#### *\$tls\_in\_bits*

Contains an approximation of the TLS cipher's bit-strength on the inbound connection; the meaning of this depends upon the TLS implementation used. If TLS has not been negotiated, the value will be 0. The value of this is automatically fed into the Cyrus SASL authenticator when acting as a server, to specify the "external SSF" (a SASL term).

The deprecated *\$tls\_bits* variable refers to the inbound side except when used in the context of an outbound SMTP delivery, when it refers to the outbound.

#### *\$tls\_out\_bits*

Contains an approximation of the TLS cipher's bit-strength on an outbound SMTP connection; the meaning of this depends upon the TLS implementation used. If TLS has not been negotiated, the value will be 0.

#### *\$tls\_in\_ourcert*

This variable refers to the certificate presented to the peer of an inbound connection when the message was received. It is only useful as the argument of a **certextract** expansion item, **md5**, **sha1** or **sha256** operator, or a **def** condition.

**Note:** Under current versions of OpenSSL, when a list of more than one file is used for **tls\_certificate**, this variable is not reliable.

#### *\$tls\_in\_peercert*

This variable refers to the certificate presented by the peer of an inbound connection when the message was received. It is only useful as the argument of a **certextract** expansion item, **md5**, **sha1** or **sha256** operator, or a **def** condition. If certificate verification fails it may refer to a failing chain element which is not the leaf.

#### *\$tls\_out\_ourcert*

This variable refers to the certificate presented to the peer of an outbound connection. It is only useful as the argument of a **certextract** expansion item, **md5**, **sha1** or **sha256** operator, or a **def** condition.

#### *\$tls\_out\_peercert*

This variable refers to the certificate presented by the peer of an outbound connection. It is only useful as the argument of a **certextract** expansion item, **md5**, **sha1** or **sha256** operator, or a **def** condition. If certificate verification fails it may refer to a failing chain element which is not the leaf.

#### *\$tls\_in\_certificate\_verified*

This variable is set to “1” if a TLS certificate was verified when the message was received, and “0” otherwise.

The deprecated *\$tls\_certificate\_verified* variable refers to the inbound side except when used in the context of an outbound SMTP delivery, when it refers to the outbound.

#### *\$tls\_out\_certificate\_verified*

This variable is set to “1” if a TLS certificate was verified when an outbound SMTP connection was made, and “0” otherwise.

#### *\$tls\_in\_cipher*

When a message is received from a remote host over an encrypted SMTP connection, this variable is set to the cipher suite that was negotiated, for example DES-CBC3-SHA. In other circumstances, in particular, for message received over unencrypted connections, the variable is empty. Testing *\$tls\_in\_cipher* for emptiness is one way of distinguishing between encrypted and non-encrypted connections during ACL processing.

The deprecated *\$tls\_cipher* variable is the same as *\$tls\_in\_cipher* during message reception, but in the context of an outward SMTP delivery taking place via the *smtp* transport becomes the same as *\$tls\_out\_cipher*.

#### *\$tls\_out\_cipher*

This variable is cleared before any outgoing SMTP connection is made, and then set to the outgoing cipher suite if one is negotiated. See chapter 42 for details of TLS support and chapter 30 for details of the *smtp* transport.

#### *\$tls\_in\_ocsp*

When a message is received from a remote client connection the result of any OCSP request from the client is encoded in this variable:

- 0 OCSP proof was not requested (default value)
- 1 No response to request
- 2 Response not verified
- 3 Verification failed
- 4 Verification succeeded

#### *\$tls\_out\_ocsp*

When a message is sent to a remote host connection the result of any OCSP request made is encoded in this variable. See *\$tls\_in\_ocsp* for values.

#### *\$tls\_in\_peerdn*

When a message is received from a remote host over an encrypted SMTP connection, and Exim is configured to request a certificate from the client, the value of the Distinguished Name of the certificate is made available in the *\$tls\_in\_peerdn* during subsequent processing. If certificate verification fails it may refer to a failing chain element which is not the leaf.

The deprecated *\$tls\_peerdn* variable refers to the inbound side except when used in the context of an outbound SMTP delivery, when it refers to the outbound.

#### *\$tls\_out\_peerdn*

When a message is being delivered to a remote host over an encrypted SMTP connection, and Exim is configured to request a certificate from the server, the value of the Distinguished Name of the certificate is made available in the *\$tls\_out\_peerdn* during subsequent processing. If certificate verification fails it may refer to a failing chain element which is not the leaf.

#### *\$tls\_in\_sni*

When a TLS session is being established, if the client sends the Server Name Indication extension, the value will be placed in this variable. If the variable appears in **tls\_certificate** then this option and some others, described in 42.10, will be re-expanded early in the TLS session, to permit a different certificate to be presented (and optionally a different key to be used) to the client, based upon the value of the SNI extension.

The deprecated *\$tls\_sni* variable refers to the inbound side except when used in the context of an outbound SMTP delivery, when it refers to the outbound.

#### *\$tls\_out\_sni*

During outbound SMTP deliveries, this variable reflects the value of the **tls\_sni** option on the transport.

#### *\$tod\_bsdinbox*

The time of day and the date, in the format required for BSD-style mailbox files, for example: Thu Oct 17 17:14:09 1995.

#### *\$tod\_epoch*

The time and date as a number of seconds since the start of the Unix epoch.

#### *\$tod\_epoch\_l*

The time and date as a number of microseconds since the start of the Unix epoch.

#### *\$tod\_full*

A full version of the time and date, for example: Wed, 16 Oct 1995 09:51:40 +0100. The timezone is always given as a numerical offset from UTC, with positive values used for timezones that are ahead (east) of UTC, and negative values for those that are behind (west).

#### *\$tod\_log*

The time and date in the format used for writing Exim's log files, for example: 1995-10-12 15:32:29, but without a timezone.

#### *\$tod\_logfile*

This variable contains the date in the format `yyymmdd`. This is the format that is used for timestamping log files when **log\_file\_path** contains the `%D` flag.

#### *\$tod\_zone*

This variable contains the numerical value of the local timezone, for example: -0500.

#### *\$tod\_zulu*

This variable contains the UTC date and time in "Zulu" format, as specified by ISO 8601, for example: 20030221154023Z.

#### *\$transport\_name*

During the running of a transport, this variable contains its name.

#### *\$value*

This variable contains the result of an expansion lookup, extraction operation, or external command, as described above. It is also used during a **reduce** expansion.

#### *\$verify\_mode*

While a router or transport is being run in verify mode or for cutthrough delivery, contains "S" for sender-verification or "R" for recipient-verification. Otherwise, empty.

#### *\$version\_number*

The version number of Exim.

#### *\$warn\_message\_delay*

This variable is set only during the creation of a message warning about a delivery delay. Details of its use are explained in section 49.2.

#### *\$warn\_message\_recipients*

This variable is set only during the creation of a message warning about a delivery delay. Details of its use are explained in section 49.2.



## 12. Embedded Perl

Exim can be built to include an embedded Perl interpreter. When this is done, Perl subroutines can be called as part of the string expansion process. To make use of the Perl support, you need version 5.004 or later of Perl installed on your system. To include the embedded interpreter in the Exim binary, include the line

```
EXIM_PERL = perl.o
```

in your *Local/Makefile* and then build Exim in the normal way.

### 12.1 Setting up so Perl can be used

Access to Perl subroutines is via a global configuration option called **perl\_startup** and an expansion string operator **\${perl ...}**. If there is no **perl\_startup** option in the Exim configuration file then no Perl interpreter is started and there is almost no overhead for Exim (since none of the Perl library will be paged in unless used). If there is a **perl\_startup** option then the associated value is taken to be Perl code which is executed in a newly created Perl interpreter.

The value of **perl\_startup** is not expanded in the Exim sense, so you do not need backslashes before any characters to escape special meanings. The option should usually be something like

```
perl_startup = do '/etc/exim.pl'
```

where */etc/exim.pl* is Perl code which defines any subroutines you want to use from Exim. Exim can be configured either to start up a Perl interpreter as soon as it is entered, or to wait until the first time it is needed. Starting the interpreter at the beginning ensures that it is done while Exim still has its *setuid* privilege, but can impose an unnecessary overhead if Perl is not in fact used in a particular run. Also, note that this does not mean that Exim is necessarily running as root when Perl is called at a later time. By default, the interpreter is started only when it is needed, but this can be changed in two ways:

- Setting **perl\_at\_start** (a boolean option) in the configuration requests a startup when Exim is entered.
- The command line option **-ps** also requests a startup when Exim is entered, overriding the setting of **perl\_at\_start**.

There is also a command line option **-pd** (for delay) which suppresses the initial startup, even if **perl\_at\_start** is set.

- To provide more security executing Perl code via the embedded Perl interpreter, the **perl\_taintmode** option can be set. This enables the taint mode of the Perl interpreter. You are encouraged to set this option to a true value. To avoid breaking existing installations, it defaults to false.

### 12.2 Calling Perl subroutines

When the configuration file includes a **perl\_startup** option you can make use of the string expansion item to call the Perl subroutines that are defined by the **perl\_startup** code. The operator is used in any of the following forms:

```
${perl{foo}}  
${perl{foo}{argument}}  
${perl{foo}{argument1}{argument2} ... }
```

which calls the subroutine **foo** with the given arguments. A maximum of eight arguments may be passed. Passing more than this results in an expansion failure with an error message of the form

```
Too many arguments passed to Perl subroutine "foo" (max is 8)
```

The return value of the Perl subroutine is evaluated in a scalar context before it is passed back to Exim to be inserted into the expanded string. If the return value is *undef*, the expansion is forced to fail in the same way as an explicit “fail” on an **if** or **lookup** item. If the subroutine aborts by obeying Perl’s **die** function, the expansion fails with the error message that was passed to **die**.

## 12.3 Calling Exim functions from Perl

Within any Perl code called from Exim, the function *Exim::expand\_string()* is available to call back into Exim's string expansion function. For example, the Perl code

```
my $lp = Exim::expand_string('$local_part');
```

makes the current Exim *\$local\_part* available in the Perl variable *\$lp*. Note those are single quotes and not double quotes to protect against *\$local\_part* being interpolated as a Perl variable.

If the string expansion is forced to fail by a “fail” item, the result of *Exim::expand\_string()* is **undef**. If there is a syntax error in the expansion string, the Perl call from the original expansion string fails with an appropriate error message, in the same way as if **die** were used.

Two other Exim functions are available for use from within Perl code. *Exim::debug\_write()* writes a string to the standard error stream if Exim's debugging is enabled. If you want a newline at the end, you must supply it. *Exim::log\_write()* writes a string to Exim's main log, adding a leading timestamp. In this case, you should not supply a terminating newline.

## 12.4 Use of standard output and error by Perl

You should not write to the standard error or output streams from within your Perl code, as it is not defined how these are set up. In versions of Exim before 4.50, it is possible for the standard output or error to refer to the SMTP connection during message reception via the daemon. Writing to this stream is certain to cause chaos. From Exim 4.50 onwards, the standard output and error streams are connected to */dev/null* in the daemon. The chaos is avoided, but the output is lost.

The Perl **warn** statement writes to the standard error stream by default. Calls to **warn** may be embedded in Perl modules that you use, but over which you have no control. When Exim starts up the Perl interpreter, it arranges for output from the **warn** statement to be written to the Exim main log. You can change this by including appropriate Perl magic somewhere in your Perl code. For example, to discard **warn** output completely, you need this:

```
$SIG{__WARN__} = sub { };
```

Whenever a **warn** is obeyed, the anonymous subroutine is called. In this example, the code for the subroutine is empty, so it does nothing, but you can include any Perl code that you like. The text of the **warn** message is passed as the first subroutine argument.

## 13. Starting the daemon and the use of network interfaces

A host that is connected to a TCP/IP network may have one or more physical hardware network interfaces. Each of these interfaces may be configured as one or more “logical” interfaces, which are the entities that a program actually works with. Each of these logical interfaces is associated with an IP address. In addition, TCP/IP software supports “loopback” interfaces (127.0.0.1 in IPv4 and ::1 in IPv6), which do not use any physical hardware. Exim requires knowledge about the host’s interfaces for use in three different circumstances:

- (1) When a listening daemon is started, Exim needs to know which interfaces and ports to listen on.
- (2) When Exim is routing an address, it needs to know which IP addresses are associated with local interfaces. This is required for the correct processing of MX lists by removing the local host and others with the same or higher priority values. Also, Exim needs to detect cases when an address is routed to an IP address that in fact belongs to the local host. Unless the **self** router option or the **allow\_localhost** option of the smtp transport is set (as appropriate), this is treated as an error situation.
- (3) When Exim connects to a remote host, it may need to know which interface to use for the outgoing connection.

Exim’s default behaviour is likely to be appropriate in the vast majority of cases. If your host has only one interface, and you want all its IP addresses to be treated in the same way, and you are using only the standard SMTP port, you should not need to take any special action. The rest of this chapter does not apply to you.

In a more complicated situation you may want to listen only on certain interfaces, or on different ports, and for this reason there are a number of options that can be used to influence Exim’s behaviour. The rest of this chapter describes how they operate.

When a message is received over TCP/IP, the interface and port that were actually used are set in *\$received\_ip\_address* and *\$received\_port*.

### 13.1 Starting a listening daemon

When a listening daemon is started (by means of the **-bd** command line option), the interfaces and ports on which it listens are controlled by the following options:

- **daemon\_smtp\_ports** contains a list of default ports or service names. (For backward compatibility, this option can also be specified in the singular.)
- **local\_interfaces** contains list of interface IP addresses on which to listen. Each item may optionally also specify a port.

The default list separator in both cases is a colon, but this can be changed as described in section 6.20. When IPv6 addresses are involved, it is usually best to change the separator to avoid having to double all the colons. For example:

```
local_interfaces = <; 127.0.0.1 ; \
                    192.168.23.65 ; \
                    ::1 ; \
                    3ffe:ffff:836f::fe86:a061
```

There are two different formats for specifying a port along with an IP address in **local\_interfaces**:

- (1) The port is added onto the address with a dot separator. For example, to listen on port 1234 on two different IP addresses:

```
local_interfaces = <; 192.168.23.65.1234 ; \
                    3ffe:ffff:836f::fe86:a061.1234
```

- (2) The IP address is enclosed in square brackets, and the port is added with a colon separator, for example:

```
local_interfaces = <; [192.168.23.65]:1234 ; \  
[3ffe:ffff:836f::fe86:a061]:1234
```

When a port is not specified, the value of **daemon\_smtp\_ports** is used. The default setting contains just one port:

```
daemon_smtp_ports = smtp
```

If more than one port is listed, each interface that does not have its own port specified listens on all of them. Ports that are listed in **daemon\_smtp\_ports** can be identified either by name (defined in */etc/services*) or by number. However, when ports are given with individual IP addresses in **local\_interfaces**, only numbers (not names) can be used.

## 13.2 Special IP listening addresses

The addresses 0.0.0.0 and ::0 are treated specially. They are interpreted as “all IPv4 interfaces” and “all IPv6 interfaces”, respectively. In each case, Exim tells the TCP/IP stack to “listen on all IPv $x$  interfaces” instead of setting up separate listening sockets for each interface. The default value of **local\_interfaces** is

```
local_interfaces = 0.0.0.0
```

when Exim is built without IPv6 support; otherwise it is:

```
local_interfaces = <; ::0 ; 0.0.0.0
```

Thus, by default, Exim listens on all available interfaces, on the SMTP port.

## 13.3 Overriding local\_interfaces and daemon\_smtp\_ports

The **-oX** command line option can be used to override the values of **daemon\_smtp\_ports** and/or **local\_interfaces** for a particular daemon instance. Another way of doing this would be to use macros and the **-D** option. However, **-oX** can be used by any admin user, whereas modification of the runtime configuration by **-D** is allowed only when the caller is root or exim.

The value of **-oX** is a list of items. The default colon separator can be changed in the usual way if required. If there are any items that do not contain dots or colons (that is, are not IP addresses), the value of **daemon\_smtp\_ports** is replaced by the list of those items. If there are any items that do contain dots or colons, the value of **local\_interfaces** is replaced by those items. Thus, for example,

```
-oX 1225
```

overrides **daemon\_smtp\_ports**, but leaves **local\_interfaces** unchanged, whereas

```
-oX 192.168.34.5.1125
```

overrides **local\_interfaces**, leaving **daemon\_smtp\_ports** unchanged. (However, since **local\_interfaces** now contains no items without ports, the value of **daemon\_smtp\_ports** is no longer relevant in this example.)

## 13.4 Support for the obsolete SSMTP (or SMTPS) protocol

Exim supports the obsolete SSMTP protocol (also known as SMTPS) that was used before the STARTTLS command was standardized for SMTP. Some legacy clients still use this protocol. If the **tls\_on\_connect\_ports** option is set to a list of port numbers or service names, connections to those ports must use SSMTP. The most common use of this option is expected to be

```
tls_on_connect_ports = 465
```

because 465 is the usual port number used by the legacy clients. There is also a command line option **-tls-on-connect**, which forces all ports to behave in this way when a daemon is started.

**Warning:** Setting **tls\_on\_connect\_ports** does not of itself cause the daemon to listen on those ports. You must still specify them in **daemon\_smtp\_ports**, **local\_interfaces**, or the **-oX** option. (This is

because **tls\_on\_connect\_ports** applies to **inetd** connections as well as to connections via the daemon.)

## 13.5 IPv6 address scopes

IPv6 addresses have “scopes”, and a host with multiple hardware interfaces can, in principle, have the same link-local IPv6 address on different interfaces. Thus, additional information is needed, over and above the IP address, to distinguish individual interfaces. A convention of using a percent sign followed by something (often the interface name) has been adopted in some cases, leading to addresses like this:

```
fe80::202:b3ff:fe03:45c1%eth0
```

To accommodate this usage, a percent sign followed by an arbitrary string is allowed at the end of an IPv6 address. By default, Exim calls *getaddrinfo()* to convert a textual IPv6 address for actual use. This function recognizes the percent convention in operating systems that support it, and it processes the address appropriately. Unfortunately, some older libraries have problems with *getaddrinfo()*. If

```
IPV6_USE_INET_PTON=yes
```

is set in *Local/Makefile* (or an OS-dependent Makefile) when Exim is built, Exim uses *inet\_pton()* to convert a textual IPv6 address for actual use, instead of *getaddrinfo()*. (Before version 4.14, it always used this function.) Of course, this means that the additional functionality of *getaddrinfo()* – recognizing scoped addresses – is lost.

## 13.6 Disabling IPv6

Sometimes it happens that an Exim binary that was compiled with IPv6 support is run on a host whose kernel does not support IPv6. The binary will fall back to using IPv4, but it may waste resources looking up AAAA records, and trying to connect to IPv6 addresses, causing delays to mail delivery. If you set the **disable\_ipv6** option true, even if the Exim binary has IPv6 support, no IPv6 activities take place. AAAA records are never looked up, and any IPv6 addresses that are listed in **local\_interfaces**, data for the *manualroute* router, etc. are ignored. If IP literals are enabled, the *ipliteral* router declines to handle IPv6 literal addresses.

On the other hand, when IPv6 is in use, there may be times when you want to disable it for certain hosts or domains. You can use the **dns\_ipv4\_lookup** option to globally suppress the lookup of AAAA records for specified domains, and you can use the **ignore\_target\_hosts** generic router option to ignore IPv6 addresses in an individual router.

## 13.7 Examples of starting a listening daemon

The default case in an IPv6 environment is

```
daemon_smtp_ports = smtp
local_interfaces = <; ::0 ; 0.0.0.0
```

This specifies listening on the smtp port on all IPv6 and IPv4 interfaces. Either one or two sockets may be used, depending on the characteristics of the TCP/IP stack. (This is complicated and messy; for more information, read the comments in the *daemon.c* source file.)

To specify listening on ports 25 and 26 on all interfaces:

```
daemon_smtp_ports = 25 : 26
```

(leaving **local\_interfaces** at the default setting) or, more explicitly:

```
local_interfaces = <; ::0.25      ; ::0.26 \
                  0.0.0.0.25 ; 0.0.0.0.26
```

To listen on the default port on all IPv4 interfaces, and on port 26 on the IPv4 loopback address only:

```
local_interfaces = 0.0.0.0 : 127.0.0.1.26
```

To specify listening on the default port on specific interfaces only:

```
local_interfaces = 10.0.0.67 : 192.168.34.67
```

**Warning:** Such a setting excludes listening on the loopback interfaces.

## 13.8 Recognizing the local host

The **local\_interfaces** option is also used when Exim needs to determine whether or not an IP address refers to the local host. That is, the IP addresses of all the interfaces on which a daemon is listening are always treated as local.

For this usage, port numbers in **local\_interfaces** are ignored. If either of the items 0.0.0.0 or ::0 are encountered, Exim gets a complete list of available interfaces from the operating system, and extracts the relevant (that is, IPv4 or IPv6) addresses to use for checking.

Some systems set up large numbers of virtual interfaces in order to provide many virtual web servers. In this situation, you may want to listen for email on only a few of the available interfaces, but nevertheless treat all interfaces as local when routing. You can do this by setting **extra\_local\_interfaces** to a list of IP addresses, possibly including the “all” wildcard values. These addresses are recognized as local, but are not used for listening. Consider this example:

```
local_interfaces = <; 127.0.0.1 ; ::1 ; \  
                  192.168.53.235 ; \  
                  3ffe:2101:12:1:a00:20ff:fe86:a061  
  
extra_local_interfaces = <; ::0 ; 0.0.0.0
```

The daemon listens on the loopback interfaces and just one IPv4 and one IPv6 address, but all available interface addresses are treated as local when Exim is routing.

In some environments the local host name may be in an MX list, but with an IP address that is not assigned to any local interface. In other cases it may be desirable to treat other host names as if they referred to the local host. Both these cases can be handled by setting the **hosts\_treat\_as\_local** option. This contains host names rather than IP addresses. When a host is referenced during routing, either via an MX record or directly, it is treated as the local host if its name matches **hosts\_treat\_as\_local**, or if any of its IP addresses match **local\_interfaces** or **extra\_local\_interfaces**.

## 13.9 Delivering to a remote host

Delivery to a remote host is handled by the smtp transport. By default, it allows the system’s TCP/IP functions to choose which interface to use (if there is more than one) when connecting to a remote host. However, the **interface** option can be set to specify which interface is used. See the description of the smtp transport in chapter 30 for more details.

## 14. Main configuration

The first part of the run time configuration file contains three types of item:

- Macro definitions: These lines start with an upper case letter. See section 6.4 for details of macro processing.
- Named list definitions: These lines start with one of the words “domainlist”, “hostlist”, “addresslist”, or “localpartlist”. Their use is described in section 10.5.
- Main configuration settings: Each setting occupies one line of the file (with possible continuations). If any setting is preceded by the word “hide”, the **-bP** command line option displays its value to admin users only. See section 6.11 for a description of the syntax of these option settings.

This chapter specifies all the main configuration options, along with their types and default values. For ease of finding a particular option, they appear in alphabetical order in section 14.23 below. However, because there are now so many options, they are first listed briefly in functional groups, as an aid to finding the name of the option you are looking for. Some options are listed in more than one group.

### 14.1 Miscellaneous

<b>bi_command</b>	to run for <b>-bi</b> command line option
<b>debug_store</b>	do extra internal checks
<b>disable_ipv6</b>	do no IPv6 processing
<b>keep_malformed</b>	for broken files – should not happen
<b>localhost_number</b>	for unique message ids in clusters
<b>message_body_newlines</b>	retain newlines in <i>\$message_body</i>
<b>message_body_visible</b>	how much to show in <i>\$message_body</i>
<b>mua_wrapper</b>	run in “MUA wrapper” mode
<b>print_tophitchars</b>	top-bit characters are printing
<b>spool_wireformat</b>	use wire-format spool data files when possible
<b>timezone</b>	force time zone

### 14.2 Exim parameters

<b>exim_group</b>	override compiled-in value
<b>exim_path</b>	override compiled-in value
<b>exim_user</b>	override compiled-in value
<b>primary_hostname</b>	default from <i>uname()</i>
<b>split_spool_directory</b>	use multiple directories
<b>spool_directory</b>	override compiled-in value

### 14.3 Privilege controls

<b>admin_groups</b>	groups that are Exim admin users
<b>commandline_checks_require_admin</b>	require admin for various checks
<b>deliver_drop_privilege</b>	drop root for delivery processes
<b>local_from_check</b>	insert <i>Sender:</i> if necessary
<b>local_from_prefix</b>	for testing <i>From:</i> for local sender
<b>local_from_suffix</b>	for testing <i>From:</i> for local sender
<b>local_sender_retain</b>	keep <i>Sender:</i> from untrusted user
<b>never_users</b>	do not run deliveries as these
<b>prod_requires_admin</b>	forced delivery requires admin user
<b>queue_list_requires_admin</b>	queue listing requires admin user
<b>trusted_groups</b>	groups that are trusted
<b>trusted_users</b>	users that are trusted

## 14.4 Logging

<code>event_action</code>	custom logging
<code>hosts_connection_nolog</code>	exemption from connect logging
<code>log_file_path</code>	override compiled-in value
<code>log_selector</code>	set/unset optional logging
<code>log_timezone</code>	add timezone to log lines
<code>message_logs</code>	create per-message logs
<code>preserve_message_logs</code>	after message completion
<code>process_log_path</code>	for SIGUSR1 and <i>exiwhat</i>
<code>slow_lookup_log</code>	control logging of slow DNS lookups
<code>syslog_duplication</code>	controls duplicate log lines on syslog
<code>syslog_facility</code>	set syslog “facility” field
<code>syslog_pid</code>	pid in syslog lines
<code>syslog_processname</code>	set syslog “ident” field
<code>syslog_timestamp</code>	timestamp syslog lines
<code>write_rejectlog</code>	control use of message log

## 14.5 Frozen messages

<code>auto_thaw</code>	sets time for retrying frozen messages
<code>freeze_tell</code>	send message when freezing
<code>move_frozen_messages</code>	to another directory
<code>timeout_frozen_after</code>	keep frozen messages only so long

## 14.6 Data lookups

<code>ibase_servers</code>	InterBase servers
<code>ldap_ca_cert_dir</code>	dir of CA certs to verify LDAP server’s
<code>ldap_ca_cert_file</code>	file of CA certs to verify LDAP server’s
<code>ldap_cert_file</code>	client cert file for LDAP
<code>ldap_cert_key</code>	client key file for LDAP
<code>ldap_cipher_suite</code>	TLS negotiation preference control
<code>ldap_default_servers</code>	used if no server in query
<code>ldap_require_cert</code>	action to take without LDAP server cert
<code>ldap_start_tls</code>	require TLS within LDAP
<code>ldap_version</code>	set protocol version
<code>lookup_open_max</code>	lookup files held open
<code>mysql_servers</code>	default MySQL servers
<code>oracle_servers</code>	Oracle servers
<code>pgsql_servers</code>	default PostgreSQL servers
<code>sqlite_lock_timeout</code>	as it says

## 14.7 Message ids

<code>message_id_header_domain</code>	used to build <i>Message-ID</i> : header
<code>message_id_header_text</code>	ditto

## 14.8 Embedded Perl Startup

<code>perl_at_start</code>	always start the interpreter
<code>perl_startup</code>	code to obey when starting Perl
<code>perl_taintmode</code>	enable taint mode in Perl

## 14.9 Daemon

<code>daemon_smtp_ports</code>	default ports
<code>daemon_startup_retries</code>	number of times to retry



daemon\_startup\_sleep  
extra\_local\_interfaces  
local\_interfaces  
pid\_file\_path  
queue\_run\_max

time to sleep between tries  
not necessarily listened on  
on which to listen, with optional ports  
override compiled-in value  
maximum simultaneous queue runners

## 14.10 Resource control

check\_log\_inodes  
check\_log\_space  
check\_spool\_inodes  
check\_spool\_space  
deliver\_queue\_load\_max  
queue\_only\_load  
queue\_only\_load\_latch  
queue\_run\_max  
remote\_max\_parallel  
smtp\_accept\_max  
smtp\_accept\_max\_nonmail  
smtp\_accept\_max\_nonmail\_hosts  
smtp\_accept\_max\_per\_connection  
smtp\_accept\_max\_per\_host  
smtp\_accept\_queue  
smtp\_accept\_queue\_per\_connection  
smtp\_accept\_reserve  
smtp\_check\_spool\_space  
smtp\_connect\_backlog  
smtp\_load\_reserve  
smtp\_reserve\_hosts

before accepting a message  
before accepting a message  
before accepting a message  
before accepting a message  
no queue deliveries if load high  
queue incoming if load high  
don't re-evaluate load for each message  
maximum simultaneous queue runners  
parallel SMTP delivery per message  
simultaneous incoming connections  
non-mail commands  
hosts to which the limit applies  
messages per connection  
connections from one host  
queue mail if more connections  
queue if more messages per connection  
only reserve hosts if more connections  
from SIZE on MAIL command  
passed to TCP/IP stack  
SMTP from reserved hosts if load high  
these are the reserve hosts

## 14.11 Policy controls

acl\_not\_smtp  
acl\_not\_smtp\_mime  
acl\_not\_smtp\_start  
acl\_smtp\_auth  
acl\_smtp\_connect  
acl\_smtp\_data  
acl\_smtp\_data\_prdr  
acl\_smtp\_dkim  
acl\_smtp\_etrn  
acl\_smtp\_expn  
acl\_smtp\_helo  
acl\_smtp\_mail  
acl\_smtp\_mailauth  
acl\_smtp\_mime  
acl\_smtp\_notquit  
acl\_smtp\_predata  
acl\_smtp\_quit  
acl\_smtp\_rcpt  
acl\_smtp\_starttls  
acl\_smtp\_vrfy  
av\_scanner  
check\_rfc2047\_length  
dns\_csa\_search\_limit  
dns\_csa\_use\_reverse  
header\_maxsize  
header\_line\_maxsize

ACL for non-SMTP messages  
ACL for non-SMTP MIME parts  
ACL for start of non-SMTP message  
ACL for AUTH  
ACL for connection  
ACL for DATA  
ACL for DATA, per-recipient  
ACL for DKIM verification  
ACL for ETRN  
ACL for EXPN  
ACL for EHLO or HELO  
ACL for MAIL  
ACL for AUTH on MAIL command  
ACL for MIME parts  
ACL for non-QUIT terminations  
ACL for start of data  
ACL for QUIT  
ACL for RCPT  
ACL for STARTTLS  
ACL for VRFY  
specify virus scanner  
check length of RFC 2047 “encoded words”  
control CSA parent search depth  
en/disable CSA IP reverse search  
total size of message header  
individual header line limit

<b>helo_accept_junk_hosts</b>	allow syntactic junk from these hosts
<b>helo_allow_chars</b>	allow illegal chars in HELO names
<b>helo_lookup_domains</b>	lookup hostname for these HELO names
<b>helo_try_verify_hosts</b>	HELO soft-checked for these hosts
<b>helo_verify_hosts</b>	HELO hard-checked for these hosts
<b>host_lookup</b>	host name looked up for these hosts
<b>host_lookup_order</b>	order of DNS and local name lookups
<b>hosts_proxy</b>	use proxy protocol for these hosts
<b>host_reject_connection</b>	reject connection from these hosts
<b>hosts_treat_as_local</b>	useful in some cluster configurations
<b>local_scan_timeout</b>	timeout for <i>local_scan()</i>
<b>message_size_limit</b>	for all messages
<b>percent_hack_domains</b>	recognize %-hack for these domains
<b>spamd_address</b>	set interface to SpamAssassin
<b>strict_acl_vars</b>	object to unset ACL variables

## 14.12 Callout cache

<b>callout_domain_negative_expire</b>	timeout for negative domain cache item
<b>callout_domain_positive_expire</b>	timeout for positive domain cache item
<b>callout_negative_expire</b>	timeout for negative address cache item
<b>callout_positive_expire</b>	timeout for positive address cache item
<b>callout_random_local_part</b>	string to use for “random” testing

## 14.13 TLS

<b>gnutls_compat_mode</b>	use GnuTLS compatibility mode
<b>gnutls_allow_auto_pkcs11</b>	allow GnuTLS to autoload PKCS11 modules
<b>openssl_options</b>	adjust OpenSSL compatibility options
<b>tls_advertise_hosts</b>	advertise TLS to these hosts
<b>tls_certificate</b>	location of server certificate
<b>tls_crl</b>	certificate revocation list
<b>tls_dh_max_bits</b>	clamp D-H bit count suggestion
<b>tls_dhparam</b>	DH parameters for server
<b>tls_eccurve</b>	EC curve selection for server
<b>tls_ocsp_file</b>	location of server certificate status proof
<b>tls_on_connect_ports</b>	specify SSMTP (SMTPS) ports
<b>tls_privatekey</b>	location of server private key
<b>tls_remember_esmtp</b>	don’t reset after starting TLS
<b>tls_require_ciphers</b>	specify acceptable ciphers
<b>tls_try_verify_hosts</b>	try to verify client certificate
<b>tls_verify_certificates</b>	expected client certificates
<b>tls_verify_hosts</b>	insist on client certificate verify

## 14.14 Local user handling

<b>finduser_retries</b>	useful in NIS environments
<b>gecos_name</b>	used when creating <i>Sender</i> :
<b>gecos_pattern</b>	ditto
<b>max_username_length</b>	for systems that truncate
<b>unknown_login</b>	used when no login name found
<b>unknown_username</b>	ditto
<b>uucp_from_pattern</b>	for recognizing “From ” lines
<b>uucp_from_sender</b>	ditto

## 14.15 All incoming messages (SMTP and non-SMTP)

<code>header_maxsize</code>	total size of message header
<code>header_line_maxsize</code>	individual header line limit
<code>message_size_limit</code>	applies to all messages
<code>percent_hack_domains</code>	recognize %-hack for these domains
<code>received_header_text</code>	expanded to make <i>Received:</i>
<code>received_headers_max</code>	for mail loop detection
<code>recipients_max</code>	limit per message
<code>recipients_max_reject</code>	permanently reject excess recipients

## 14.16 Non-SMTP incoming messages

<code>receive_timeout</code>	for non-SMTP messages
------------------------------	-----------------------

## 14.17 Incoming SMTP messages

See also the *Policy controls* section above.

<code>dkim_verify_signers</code>	DKIM domain for which DKIM ACL is run
<code>host_lookup</code>	host name looked up for these hosts
<code>host_lookup_order</code>	order of DNS and local name lookups
<code>recipient_unqualified_hosts</code>	may send unqualified recipients
<code>rfc1413_hosts</code>	make ident calls to these hosts
<code>rfc1413_query_timeout</code>	zero disables ident calls
<code>sender_unqualified_hosts</code>	may send unqualified senders
<code>smtp_accept_keepalive</code>	some TCP/IP magic
<code>smtp_accept_max</code>	simultaneous incoming connections
<code>smtp_accept_max_nonmail</code>	non-mail commands
<code>smtp_accept_max_nonmail_hosts</code>	hosts to which the limit applies
<code>smtp_accept_max_per_connection</code>	messages per connection
<code>smtp_accept_max_per_host</code>	connections from one host
<code>smtp_accept_queue</code>	queue mail if more connections
<code>smtp_accept_queue_per_connection</code>	queue if more messages per connection
<code>smtp_accept_reserve</code>	only reserve hosts if more connections
<code>smtp_active_hostname</code>	host name to use in messages
<code>smtp_banner</code>	text for welcome banner
<code>smtp_check_spool_space</code>	from SIZE on MAIL command
<code>smtp_connect_backlog</code>	passed to TCP/IP stack
<code>smtp_enforce_sync</code>	of SMTP command/responses
<code>smtp_etrn_command</code>	what to run for ETRN
<code>smtp_etrn_serialize</code>	only one at once
<code>smtp_load_reserve</code>	only reserve hosts if this load
<code>smtp_max_unknown_commands</code>	before dropping connection
<code>smtp_ratelimit_hosts</code>	apply ratelimiting to these hosts
<code>smtp_ratelimit_mail</code>	ratelimit for MAIL commands
<code>smtp_ratelimit_rcpt</code>	ratelimit for RCPT commands
<code>smtp_receive_timeout</code>	per command or data line
<code>smtp_reserve_hosts</code>	these are the reserve hosts
<code>smtp_return_error_details</code>	give detail on rejections

## 14.18 SMTP extensions

<code>accept_8bitmime</code>	advertise 8BITMIME
<code>auth_advertise_hosts</code>	advertise AUTH to these hosts
<code>chunking_advertise_hosts</code>	advertise CHUNKING to these hosts
<code>dsn_advertise_hosts</code>	advertise DSN extensions to these hosts
<code>ignore_fromline_hosts</code>	allow "From " from these hosts
<code>ignore_fromline_local</code>	allow "From " from local SMTP

pipelining\_advertise\_hosts  
prdr\_enable  
smtputf8\_advertise\_hosts  
tls\_advertise\_hosts

advertise pipelining to these hosts  
advertise PRDR to all hosts  
advertise SMTPUTF8 to these hosts  
advertise TLS to these hosts

## 14.19 Processing messages

allow\_domain\_literals  
allow\_mx\_to\_ip  
allow\_utf8\_domains  
check\_rfc2047\_length  
delivery\_date\_remove  
envelope\_to\_remove  
extract\_addresses\_remove\_arguments  
headers\_charset  
qualify\_domain  
qualify\_recipient  
return\_path\_remove  
strip\_excess\_angle\_brackets  
strip\_trailing\_dot  
untrusted\_set\_sender

recognize domain literal syntax  
allow MX to point to IP address  
in addresses  
check length of RFC 2047 “encoded words”  
from incoming messages  
from incoming messages  
affects **-t** processing  
default for translations  
default for senders  
default for recipients  
from incoming messages  
in addresses  
at end of addresses  
untrusted can set envelope sender

## 14.20 System filter

system\_filter  
system\_filter\_directory\_transport  
system\_filter\_file\_transport  
system\_filter\_group  
system\_filter\_pipe\_transport  
system\_filter\_reply\_transport  
system\_filter\_user

locate system filter  
transport for delivery to a directory  
transport for delivery to a file  
group for filter running  
transport for delivery to a pipe  
transport for autoreply delivery  
user for filter running

## 14.21 Routing and delivery

disable\_ipv6  
dns\_again\_means\_nonexist  
dns\_check\_names\_pattern  
dns\_dnssec\_ok  
dns\_ipv4\_lookup  
dns\_retrans  
dns\_retry  
dns\_trust\_aa  
dns\_use\_edns0  
hold\_domains  
local\_interfaces  
queue\_domains  
queue\_only  
queue\_only\_file  
queue\_only\_load  
queue\_only\_load\_latch  
queue\_only\_override  
queue\_run\_in\_order  
queue\_run\_max  
queue\_smtp\_domains  
remote\_max\_parallel  
remote\_sort\_domains  
retry\_data\_expire  
retry\_interval\_max

do no IPv6 processing  
for broken domains  
pre-DNS syntax check  
parameter for resolver  
only v4 lookup for these domains  
parameter for resolver  
parameter for resolver  
DNS zones trusted as authentic  
parameter for resolver  
hold delivery for these domains  
for routing checks  
no immediate delivery for these  
no immediate delivery at all  
no immediate delivery if file exists  
no immediate delivery if load is high  
don’t re-evaluate load for each message  
allow command line to override  
order of arrival  
of simultaneous queue runners  
no immediate SMTP delivery for these  
parallel SMTP delivery per message  
order of remote deliveries  
timeout for retry data  
safety net for retry rules

## 14.22 Bounce and warning messages

<b>bounce_message_file</b>	content of bounce
<b>bounce_message_text</b>	content of bounce
<b>bounce_return_body</b>	include body if returning message
<b>bounce_return_linesize_limit</b>	limit on returned message line length
<b>bounce_return_message</b>	include original message in bounce
<b>bounce_return_size_limit</b>	limit on returned message
<b>bounce_sender_authentication</b>	send authenticated sender with bounce
<b>dsn_from</b>	set <i>From:</i> contents in bounces
<b>errors_copy</b>	copy bounce messages
<b>errors_reply_to</b>	<i>Reply-to:</i> in bounces
<b>delay_warning</b>	time schedule
<b>delay_warning_condition</b>	condition for warning messages
<b>ignore_bounce_errors_after</b>	discard undeliverable bounces
<b>smtp_return_error_details</b>	give detail on rejections
<b>warn_message_file</b>	content of warning message

## 14.23 Alphabetical list of main options

Those options that undergo string expansion before use are marked with †.

<b>accept_8bitmime</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
------------------------	------------------	----------------------	----------------------

This option causes Exim to send 8BITMIME in its response to an SMTP EHLO command, and to accept the BODY= parameter on MAIL commands. However, though Exim is 8-bit clean, it is not a protocol converter, and it takes no steps to do anything special with messages received by this route.

Historically Exim kept this option off by default, but the maintainers feel that in today's Internet, this causes more problems than it solves. It now defaults to true. A more detailed analysis of the issues is provided by Dan Bernstein:

<http://cr.yp.to/smtp/8bitmime.html>

To log received 8BITMIME status use

```
log_selector = +8bitmime
```

<b>acl_not_smtp</b>	Use: <i>main</i>	Type: <i>string</i> †	Default: <i>unset</i>
---------------------	------------------	-----------------------	-----------------------

This option defines the ACL that is run when a non-SMTP message has been read and is on the point of being accepted. See chapter 43 for further details.

<b>acl_not_smtp_mime</b>	Use: <i>main</i>	Type: <i>string</i> †	Default: <i>unset</i>
--------------------------	------------------	-----------------------	-----------------------

This option defines the ACL that is run for individual MIME parts of non-SMTP messages. It operates in exactly the same way as **acl\_smtp\_mime** operates for SMTP messages.

<b>acl_not_smtp_start</b>	Use: <i>main</i>	Type: <i>string</i> †	Default: <i>unset</i>
---------------------------	------------------	-----------------------	-----------------------

This option defines the ACL that is run before Exim starts reading a non-SMTP message. See chapter 43 for further details.

<b>acl_smtp_auth</b>	Use: <i>main</i>	Type: <i>string</i> †	Default: <i>unset</i>
----------------------	------------------	-----------------------	-----------------------

This option defines the ACL that is run when an SMTP AUTH command is received. See chapter 43 for further details.

<b>acl_smtp_connect</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP connection is received. See chapter 43 for further details.

<b>acl_smtp_data</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run after an SMTP DATA command has been processed and the message itself has been received, but before the final acknowledgment is sent. See chapter 43 for further details.

<b>acl_smtp_data_prdr</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>accept</i>
---------------------------	------------------	----------------------	------------------------

This option defines the ACL that, if the PRDR feature has been negotiated, is run for each recipient after an SMTP DATA command has been processed and the message itself has been received, but before the acknowledgment is sent. See chapter 43 for further details.

<b>acl_smtp_dkim</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run for each DKIM signature (by default, or as specified in the `dkim_verify_signers` option) of a received message. See chapter 57 for further details.

<b>acl_smtp_etrn</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP ETRN command is received. See chapter 43 for further details.

<b>acl_smtp_expn</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP EXPN command is received. See chapter 43 for further details.

<b>acl_smtp_helo</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP EHLO or HELO command is received. See chapter 43 for further details.

<b>acl_smtp_mail</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP MAIL command is received. See chapter 43 for further details.

<b>acl_smtp_mailauth</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when there is an AUTH parameter on a MAIL command. See chapter 43 for details of ACLs, and chapter 33 for details of authentication.

<b>acl_smtp_mime</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option is available when Exim is built with the content-scanning extension. It defines the ACL that is run for each MIME part in a message. See section 44.4 for details.

<b>acl_smtp_notquit</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP session ends without a QUIT command being received. See chapter 43 for further details.

<b>acl_smtp_predata</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP DATA command is received, before the message itself is received. See chapter 43 for further details.

<b>acl_smtp_quit</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP QUIT command is received. See chapter 43 for further details.

<b>acl_smtp_rcpt</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP RCPT command is received. See chapter 43 for further details.

<b>acl_smtp_starttls</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP STARTTLS command is received. See chapter 43 for further details.

<b>acl_smtp_vrfy</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option defines the ACL that is run when an SMTP VRFY command is received. See chapter 43 for further details.

<b>add_environment</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>empty</i>
------------------------	------------------	--------------------------	-----------------------

This option allows to set individual environment variables that the currently linked libraries and programs in child processes use. See 29.4 for the environment of *pipe* transports.

<b>admin_groups</b>	Use: <i>main</i>	Type: <i>string list†</i>	Default: <i>unset</i>
---------------------	------------------	---------------------------	-----------------------

This option is expanded just once, at the start of Exim's processing. If the current group or any of the supplementary groups of an Exim caller is in this colon-separated list, the caller has admin privileges. If all your system programmers are in a specific group, for example, you can give them all Exim admin privileges by putting that group in **admin\_groups**. However, this does not permit them to read Exim's spool files (whose group owner is the Exim gid). To permit this, you have to add individuals to the Exim group.

<b>allow_domain_literals</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------------	------------------	----------------------	-----------------------

If this option is set, the RFC 2822 domain literal format is permitted in email addresses. The option is not set by default, because the domain literal format is not normally required these days, and few people know about it. It has, however, been exploited by mail abusers.

Unfortunately, it seems that some DNS black list maintainers are using this format to report black listing to postmasters. If you want to accept messages addressed to your hosts by IP address, you need to set **allow\_domain\_literals** true, and also to add @ [ ] to the list of local domains (defined in the

named domain list **local\_domains** in the default configuration). This “magic string” matches the domain literal form of all the local host’s IP addresses.

<b>allow_mx_to_ip</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	------------------	----------------------	-----------------------

It appears that more and more DNS zone administrators are breaking the rules and putting domain names that look like IP addresses on the right hand side of MX records. Exim follows the rules and rejects this, giving an error message that explains the misconfiguration. However, some other MTAs support this practice, so to avoid “Why can’t Exim do this?” complaints, **allow\_mx\_to\_ip** exists, in order to enable this heinous activity. It is not recommended, except when you have no other choice.

<b>allow_utf8_domains</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	------------------	----------------------	-----------------------

Lots of discussion is going on about internationalized domain names. One camp is strongly in favour of just using UTF-8 characters, and it seems that at least two other MTAs permit this. This option allows Exim users to experiment if they wish.

If it is set true, Exim’s domain parsing function allows valid UTF-8 multicharacters to appear in domain name components, in addition to letters, digits, and hyphens. However, just setting this option is not enough; if you want to look up these domain names in the DNS, you must also adjust the value of **dns\_check\_names\_pattern** to match the extended form. A suitable setting is:

```
dns_check_names_pattern = (?i)^(?>(?(1)\.|( ))[a-z0-9\xc0-\xff]\
    (?>[-a-z0-9\x80-\xff]*[a-z0-9\x80-\xbf])?)+$
```

Alternatively, you can just disable this feature by setting

```
dns_check_names_pattern =
```

That is, set the option to an empty string so that no check is done.

<b>auth_advertise_hosts</b>	Use: <i>main</i>	Type: <i>host list</i> <sup>†</sup>	Default: <i>*</i>
-----------------------------	------------------	-------------------------------------	-------------------

If any server authentication mechanisms are configured, Exim advertises them in response to an EHLO command only if the calling host matches this list. Otherwise, Exim does not advertise AUTH. Exim does not accept AUTH commands from clients to which it has not advertised the availability of AUTH. The advertising of individual authentication mechanisms can be controlled by the use of the **server\_advertise\_condition** generic authenticator option on the individual authenticators. See chapter 33 for further details.

Certain mail clients (for example, Netscape) require the user to provide a name and password for authentication if AUTH is advertised, even though it may not be needed (the host may accept messages from hosts on its local LAN without authentication, for example). The **auth\_advertise\_hosts** option can be used to make these clients more friendly by excluding them from the set of hosts to which Exim advertises AUTH.

If you want to advertise the availability of AUTH only when the connection is encrypted using TLS, you can make use of the fact that the value of this option is expanded, with a setting like this:

```
auth_advertise_hosts = ${if eq{$tls_in_cipher}{}}{*}}
```

If *\$tls\_in\_cipher* is empty, the session is not encrypted, and the result of the expansion is empty, thus matching no hosts. Otherwise, the result of the expansion is *\**, which matches all hosts.

<b>auto_thaw</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>0s</i>
------------------	------------------	-------------------	--------------------

If this option is set to a time greater than zero, a queue runner will try a new delivery attempt on any frozen message, other than a bounce message, if this much time has passed since it was frozen. This may result in the message being re-frozen if nothing has changed since the last attempt. It is a way of saying “keep on trying, even though there are big problems”.



**Note:** This is an old option, which predates **timeout\_frozen\_after** and **ignore\_bounce\_errors\_after**. It is retained for compatibility, but it is not thought to be very useful any more, and its use should probably be avoided.

<b>av_scanner</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>see below</i>
-------------------	------------------	---------------------	---------------------------

This option is available if Exim is built with the content-scanning extension. It specifies which anti-virus scanner to use. The default value is:

```
sophie:/var/run/sophie
```

If the value of **av\_scanner** starts with a dollar character, it is expanded before use. See section 44.1 for further details.

<b>bi_command</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------	------------------	---------------------	-----------------------

This option supplies the name of a command that is run when Exim is called with the **-bi** option (see chapter 5). The string value is just the command name, it is not a complete command line. If an argument is required, it must come from the **-oA** command line option.

<b>bounce_message_file</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------------	------------------	---------------------	-----------------------

This option defines a template file containing paragraphs of text to be used for constructing bounce messages. Details of the file's contents are given in chapter 49. See also **warn\_message\_file**.

<b>bounce_message_text</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------------	------------------	---------------------	-----------------------

When this option is set, its contents are included in the default bounce message immediately after "This message was created automatically by mail delivery software." It is not used if **bounce\_message\_file** is set.

<b>bounce_return_body</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------------------	------------------	----------------------	----------------------

This option controls whether the body of an incoming message is included in a bounce message when **bounce\_return\_message** is true. The default setting causes the entire message, both header and body, to be returned (subject to the value of **bounce\_return\_size\_limit**). If this option is false, only the message header is included. In the case of a non-SMTP message containing an error that is detected during reception, only those header lines preceding the point at which the error was detected are returned.

<b>bounce_return_linesize_limit</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>998</i>
-------------------------------------	------------------	----------------------	---------------------

This option sets a limit in bytes on the line length of messages that are returned to senders due to delivery problems, when **bounce\_return\_message** is true. The default value corresponds to RFC limits. If the message being returned has lines longer than this value it is treated as if the **bounce\_return\_size\_limit** (below) restriction was exceeded.

The option also applies to bounces returned when an error is detected during reception of a message. In this case lines from the original are truncated.

The option does not apply to messages generated by an *autoreply* transport.

<b>bounce_return_message</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
------------------------------	------------------	----------------------	----------------------

If this option is set false, none of the original message is included in bounce messages generated by Exim. See also **bounce\_return\_size\_limit** and **bounce\_return\_body**.

<b>bounce_return_size_limit</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>100K</i>
---------------------------------	------------------	----------------------	----------------------

This option sets a limit in bytes on the size of messages that are returned to senders as part of bounce messages when **bounce\_return\_message** is true. The limit should be less than the value of the global **message\_size\_limit** and of any **message\_size\_limit** settings on transports, to allow for the bounce text that Exim generates. If this option is set to zero there is no limit.

When the body of any message that is to be included in a bounce message is greater than the limit, it is truncated, and a comment pointing this out is added at the top. The actual cutoff may be greater than the value given, owing to the use of buffering for transferring the message in chunks (typically 8K in size). The idea is to save bandwidth on those undeliverable 15-megabyte messages.

<b>bounce_sender_authentication</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------------------------	------------------	---------------------	-----------------------

This option provides an authenticated sender address that is sent with any bounce messages generated by Exim that are sent over an authenticated SMTP connection. A typical setting might be:

```
bounce_sender_authentication = mailer-daemon@my.domain.example
```

which would cause bounce messages to be sent using the SMTP command:

```
MAIL FROM:<> AUTH=mailer-daemon@my.domain.example
```

The value of **bounce\_sender\_authentication** must always be a complete email address.

<b>callout_domain_negative_expire</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>3h</i>
---------------------------------------	------------------	-------------------	--------------------

This option specifies the expiry time for negative callout cache data for a domain. See section 43.45 for details of callout verification, and section 43.47 for details of the caching.

<b>callout_domain_positive_expire</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>7d</i>
---------------------------------------	------------------	-------------------	--------------------

This option specifies the expiry time for positive callout cache data for a domain. See section 43.45 for details of callout verification, and section 43.47 for details of the caching.

<b>callout_negative_expire</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>2h</i>
--------------------------------	------------------	-------------------	--------------------

This option specifies the expiry time for negative callout cache data for an address. See section 43.45 for details of callout verification, and section 43.47 for details of the caching.

<b>callout_positive_expire</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>24h</i>
--------------------------------	------------------	-------------------	---------------------

This option specifies the expiry time for positive callout cache data for an address. See section 43.45 for details of callout verification, and section 43.47 for details of the caching.

<b>callout_random_local_part</b>	Use: <i>main</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
----------------------------------	------------------	----------------------------------	---------------------------

This option defines the “random” local part that can be used as part of callout verification. The default value is

```
$primary_hostname-$tod_epoch-testing
```

See section 43.46 for details of how this value is used.

<b>check_log_inodes</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>100</i>
-------------------------	------------------	----------------------	---------------------

See **check\_spool\_space** below.

<b>check_log_space</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>10M</i>
------------------------	------------------	----------------------	---------------------

See **check\_spool\_space** below.

<b>check_rfc2047_length</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
-----------------------------	------------------	----------------------	----------------------

RFC 2047 defines a way of encoding non-ASCII characters in headers using a system of “encoded words”. The RFC specifies a maximum length for an encoded word; strings to be encoded that exceed this length are supposed to use multiple encoded words. By default, Exim does not recognize encoded words that exceed the maximum length. However, it seems that some software, in violation of the RFC, generates overlong encoded words. If **check\_rfc2047\_length** is set false, Exim recognizes encoded words of any length.

<b>check_spool_inodes</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>100</i>
---------------------------	------------------	----------------------	---------------------

See **check\_spool\_space** below.

<b>check_spool_space</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>10M</i>
--------------------------	------------------	----------------------	---------------------

The four **check\_...** options allow for checking of disk resources before a message is accepted.

When any of these options are nonzero, they apply to all incoming messages. If you want to apply different checks to different kinds of message, you can do so by testing the variables *\$log\_inodes*, *\$log\_space*, *\$spool\_inodes*, and *\$spool\_space* in an ACL with appropriate additional conditions.

**check\_spool\_space** and **check\_spool\_inodes** check the spool partition if either value is greater than zero, for example:

```
check_spool_space = 100M
check_spool_inodes = 100
```

The spool partition is the one that contains the directory defined by *SPOOL\_DIRECTORY* in *Local/Makefile*. It is used for holding messages in transit.

**check\_log\_space** and **check\_log\_inodes** check the partition in which log files are written if either is greater than zero. These should be set only if **log\_file\_path** and **spool\_directory** refer to different partitions.

If there is less space or fewer inodes than requested, Exim refuses to accept incoming mail. In the case of SMTP input this is done by giving a 452 temporary error response to the MAIL command. If ESMTP is in use and there was a SIZE parameter on the MAIL command, its value is added to the **check\_spool\_space** value, and the check is performed even if **check\_spool\_space** is zero, unless **no\_smtp\_check\_spool\_space** is set.

The values for **check\_spool\_space** and **check\_log\_space** are held as a number of kilobytes (though specified in bytes). If a non-multiple of 1024 is specified, it is rounded up.

For non-SMTP input and for batched SMTP input, the test is done at start-up; on failure a message is written to stderr and Exim exits with a non-zero code, as it obviously cannot send an error message of any kind.

There is a slight performance penalty for these checks. Versions of Exim preceding 4.88 had these disabled by default; high-rate installations confident they will never run out of resources may wish to deliberately disable them.

<b>chunking_advertise_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>*</i>
---------------------------------	------------------	-------------------------	-------------------

The CHUNKING extension (RFC3030) will be advertised in the EHLO message to these hosts. Hosts may use the BDAT command as an alternate to DATA.

<b>commandline_checks_require_admin</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
---	------------------	----------------------	-----------------------

This option restricts various basic checking features to require an administrative user. This affects most of the **-b\*** options, such as **-be**.

<b>debug_store</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	------------------	----------------------	-----------------------

This option, when true, enables extra checking in Exim's internal memory management. For use when a memory corruption issue is being investigated, it should normally be left as default.

<b>daemon_smtp_ports</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>smtp</i>
--------------------------	------------------	---------------------	----------------------

This option specifies one or more default SMTP ports on which the Exim daemon listens. See chapter 13 for details of how it is used. For backward compatibility, **daemon\_smtp\_port** (singular) is a synonym.

<b>daemon_startup_retries</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>9</i>
-------------------------------	------------------	----------------------	-------------------

This option, along with **daemon\_startup\_sleep**, controls the retrying done by the daemon at startup when it cannot immediately bind a listening socket (typically because the socket is already in use): **daemon\_startup\_retries** defines the number of retries after the first failure, and **daemon\_startup\_sleep** defines the length of time to wait between retries.

<b>daemon_startup_sleep</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>30s</i>
-----------------------------	------------------	-------------------	---------------------

See **daemon\_startup\_retries**.

<b>delay_warning</b>	Use: <i>main</i>	Type: <i>time list</i>	Default: <i>24h</i>
----------------------	------------------	------------------------	---------------------

When a message is delayed, Exim sends a warning message to the sender at intervals specified by this option. The data is a colon-separated list of times after which to send warning messages. If the value of the option is an empty string or a zero time, no warnings are sent. Up to 10 times may be given. If a message has been on the queue for longer than the last time, the last interval between the times is used to compute subsequent warning times. For example, with

```
delay_warning = 4h:8h:24h
```

the first message is sent after 4 hours, the second after 8 hours, and the third one after 24 hours. After that, messages are sent every 16 hours, because that is the interval between the last two times on the list. If you set just one time, it specifies the repeat interval. For example, with:

```
delay_warning = 6h
```

messages are repeated every six hours. To stop warnings after a given time, set a very large time at the end of the list. For example:

```
delay_warning = 2h:12h:99d
```

Note that the option is only evaluated at the time a delivery attempt fails, which depends on retry and queue-runner configuration. Typically retries will be configured more frequently than warning messages.

<b>delay_warning_condition</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>see below</i>
--------------------------------	------------------	----------------------	---------------------------

The string is expanded at the time a warning message might be sent. If all the deferred addresses have the same domain, it is set in *\$domain* during the expansion. Otherwise *\$domain* is empty. If the result of the expansion is a forced failure, an empty string, or a string matching any of “0”, “no” or “false” (the comparison being done caselessly) then the warning message is not sent. The default is:

```
delay_warning_condition = ${if or { \
  { !eq{$h_list-id:$h_list-post:$h_list-subscribe:}{} } \
  { match{$h_precedence:}{(?i)bulk|list|junk} } \
  { match{$h_auto-submitted:}{(?i)auto-generated|auto-replied} } \
  } {no}{yes}}
```

This suppresses the sending of warnings for messages that contain *List-ID:*, *List-Post:*, or *List-Subscribe:* headers, or have “bulk”, “list” or “junk” in a *Precedence:* header, or have “auto-generated” or “auto-replied” in an *Auto-Submitted:* header.

<b>deliver_drop_privilege</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------------	------------------	----------------------	-----------------------

If this option is set true, Exim drops its root privilege at the start of a delivery process, and runs as the Exim user throughout. This severely restricts the kinds of local delivery that are possible, but is viable in certain types of configuration. There is a discussion about the use of root privilege in chapter 55.

<b>deliver_queue_load_max</b>	Use: <i>main</i>	Type: <i>fixed-point</i>	Default: <i>unset</i>
-------------------------------	------------------	--------------------------	-----------------------

When this option is set, a queue run is abandoned if the system load average becomes greater than the value of the option. The option has no effect on ancient operating systems on which Exim cannot determine the load average. See also **queue\_only\_load** and **smtp\_load\_reserve**.

<b>delivery_date_remove</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
-----------------------------	------------------	----------------------	----------------------

Exim’s transports have an option for adding a *Delivery-date:* header to a message when it is delivered, in exactly the same way as *Return-path:* is handled. *Delivery-date:* records the actual time of delivery. Such headers should not be present in incoming messages, and this option causes them to be removed at the time the message is received, to avoid any problems that might occur when a delivered message is subsequently sent on to some other recipient.

<b>disable_fsync</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------	----------------------	-----------------------

This option is available only if Exim was built with the compile-time option `ENABLE_DISABLE_FSYNC`. When this is not set, a reference to **disable\_fsync** in a runtime configuration generates an “unknown option” error. You should not build Exim with `ENABLE_DISABLE_FSYNC` or set **disable\_fsync** unless you really, really, really understand what you are doing. *No pre-compiled distributions of Exim should ever make this option available.*

When **disable\_fsync** is set true, Exim no longer calls *fsync()* to force updated files’ data to be written to disc before continuing. Unexpected events such as crashes and power outages may cause data to be lost or scrambled. Here be Dragons. **Beware.**

<b>disable_ipv6</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------	------------------	----------------------	-----------------------

If this option is set true, even if the Exim binary has IPv6 support, no IPv6 activities take place. AAAA records are never looked up, and any IPv6 addresses that are listed in **local\_interfaces**, data for the **manualroute** router, etc. are ignored. If IP literals are enabled, the *ipliteral* router declines to handle IPv6 literal addresses.

<b>dkim_verify_signers</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>\$dkim_signers</i>
----------------------------	------------------	---------------------------	--------------------------------

This option gives a list of DKIM domains for which the DKIM ACL is run. It is expanded after the message is received; by default it runs the ACL once for each signature in the message. See chapter 57.

<b>dns_again_means_nonexist</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
---------------------------------	------------------	---------------------------	-----------------------

DNS lookups give a “try again” response for the DNS errors “non-authoritative host not found” and “SERVERFAIL”. This can cause Exim to keep trying to deliver a message, or to give repeated temporary errors to incoming mail. Sometimes the effect is caused by a badly set up name server and may persist for a long time. If a domain which exhibits this problem matches anything in **dns\_again\_means\_nonexist**, it is treated as if it did not exist. This option should be used with care. You can make it apply to reverse lookups by a setting such as this:

```
dns_again_means_nonexist = *.in-addr.arpa
```

This option applies to all DNS lookups that Exim does. It also applies when the *gethostbyname()* or *getipnodebyname()* functions give temporary errors, since these are most likely to be caused by DNS lookup problems. The *dnslookup* router has some options of its own for controlling what happens when lookups for MX or SRV records give temporary errors. These more specific options are applied after this global option.

<b>dns_check_names_pattern</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>see below</i>
--------------------------------	------------------	---------------------	---------------------------

When this option is set to a non-empty string, it causes Exim to check domain names for characters that are not allowed in host names before handing them to the DNS resolver, because some resolvers give temporary errors for names that contain unusual characters. If a domain name contains any unwanted characters, a “not found” result is forced, and the resolver is not called. The check is done by matching the domain name against a regular expression, which is the value of this option. The default pattern is

```
dns_check_names_pattern = \
  (?i)^(?>(?(1)\.|()) [^\W_] (?:[a-z0-9/-]* [^\W_] )?)+$
```

which permits only letters, digits, slashes, and hyphens in components, but they must start and end with a letter or digit. Slashes are not, in fact, permitted in host names, but they are found in certain NS records (which can be accessed in Exim by using a **dnsdb** lookup). If you set **allow\_utf8\_domains**, you must modify this pattern, or set the option to an empty string.

<b>dns_csa_search_limit</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>5</i>
-----------------------------	------------------	----------------------	-------------------

This option controls the depth of parental searching for CSA SRV records in the DNS, as described in more detail in section 43.50.

<b>dns_csa_use_reverse</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
----------------------------	------------------	----------------------	----------------------

This option controls whether or not an IP address, given as a CSA domain, is reversed and looked up in the reverse DNS, as described in more detail in section 43.50.

<b>dns_dnssec_ok</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>-1</i>
----------------------	------------------	----------------------	--------------------

If this option is set to a non-negative number then Exim will initialise the DNS resolver library to either use or not use DNSSEC, overriding the system default. A value of 0 coerces DNSSEC off, a value of 1 coerces DNSSEC on.

If the resolver library does not support DNSSEC then this option has no effect.

<b>dns_ipv4_lookup</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
------------------------	------------------	---------------------------	-----------------------

When Exim is compiled with IPv6 support and **disable\_ipv6** is not set, it looks for IPv6 address records (AAAA records) as well as IPv4 address records (A records) when trying to find IP addresses for hosts, unless the host's domain matches this list.

This is a fudge to help with name servers that give big delays or otherwise do not work for the AAAA record type. In due course, when the world's name servers have all been upgraded, there should be no need for this option.

<b>dns_retrans</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>0s</i>
--------------------	------------------	-------------------	--------------------

The options **dns\_retrans** and **dns\_retry** can be used to set the retransmission and retry parameters for DNS lookups. Values of zero (the defaults) leave the system default settings unchanged. The first value is the time between retries, and the second is the number of retries. It isn't totally clear exactly how these settings affect the total time a DNS lookup may take. I haven't found any documentation about timeouts on DNS lookups; these parameter values are available in the external resolver interface structure, but nowhere does it seem to describe how they are used or what you might want to set in them. See also the **slow\_lookup\_log** option.

<b>dns_retry</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>0</i>
------------------	------------------	----------------------	-------------------

See **dns\_retrans** above.

<b>dns_trust_aa</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
---------------------	------------------	---------------------------	-----------------------

If this option is set then lookup results marked with the AA bit (Authoritative Answer) are trusted the same way as if they were DNSSEC-verified. The authority section's name of the answer must match with this expanded domain list.

Use this option only if you talk directly to a resolver that is authoritative for some zones and does not set the AD (Authentic Data) bit in the answer. Some DNS servers may have an configuration option to mark the answers from their own zones as verified (they set the AD bit). Others do not have this option. It is considered as poor practice using a resolver that is an authoritative server for some zones.

Use this option only if you really have to (e.g. if you want to use DANE for remote delivery to a server that is listed in the DNS zones that your resolver is authoritative for).

If the DNS answer packet has the AA bit set and contains resource record in the answer section, the name of the first NS record appearing in the authority section is compared against the list. If the answer packet is authoritative but the answer section is empty, the name of the first SOA record in the authoritative section is used instead.

<b>dns_use_edns0</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>-1</i>
----------------------	------------------	----------------------	--------------------

If this option is set to a non-negative number then Exim will initialise the DNS resolver library to either use or not use EDNS0 extensions, overriding the system default. A value of 0 coerces EDNS0 off, a value of 1 coerces EDNS0 on.

If the resolver library does not support EDNS0 then this option has no effect.

OpenBSD's asr resolver routines are known to ignore the EDNS0 option; this means that DNSSEC will not work with Exim on that platform either, unless Exim is linked against an alternative DNS client library.

<b>drop_cr</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------	------------------	----------------------	-----------------------

This is an obsolete option that is now a no-op. It used to affect the way Exim handled CR and LF characters in incoming messages. What happens now is described in section 47.2.

<b>dsn_advertise_hosts</b>	Use: <i>main</i>	Type: <i>host list</i> <sup>†</sup>	Default: <i>unset</i>
----------------------------	------------------	-------------------------------------	-----------------------

DSN extensions (RFC3461) will be advertised in the EHLO message to, and accepted from, these hosts. Hosts may use the NOTIFY and ENVID options on RCPT TO commands, and RET and ORCPT options on MAIL FROM commands. A NOTIFY=SUCCESS option requests success-DSN messages. A NOTIFY= option with no argument requests that no delay or failure DSNs are sent.

<b>dsn_from</b>	Use: <i>main</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
-----------------	------------------	----------------------------------	---------------------------

This option can be used to vary the contents of *From:* header lines in bounces and other automatically generated messages ("Delivery Status Notifications" – hence the name of the option). The default setting is:

```
dsn_from = Mail Delivery System <Mailer-Daemon@$qualify_domain>
```

The value is expanded every time it is needed. If the expansion fails, a panic is logged, and the default value is used.

<b>envelope_to_remove</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------------------	------------------	----------------------	----------------------

Exim's transports have an option for adding an *Envelope-to:* header to a message when it is delivered, in exactly the same way as *Return-path:* is handled. *Envelope-to:* records the original recipient address from the message's envelope that caused the delivery to happen. Such headers should not be present in incoming messages, and this option causes them to be removed at the time the message is received, to avoid any problems that might occur when a delivered message is subsequently sent on to some other recipient.

<b>errors_copy</b>	Use: <i>main</i>	Type: <i>string list</i> <sup>†</sup>	Default: <i>unset</i>
--------------------	------------------	---------------------------------------	-----------------------

Setting this option causes Exim to send bcc copies of bounce messages that it generates to other addresses. **Note:** This does not apply to bounce messages coming from elsewhere. The value of the option is a colon-separated list of items. Each item consists of a pattern, terminated by white space, followed by a comma-separated list of email addresses. If a pattern contains spaces, it must be enclosed in double quotes.

Each pattern is processed in the same way as a single item in an address list (see section 10.19). When a pattern matches the recipient of the bounce message, the message is copied to the addresses on the list. The items are scanned in order, and once a matching one is found, no further items are examined. For example:



```

errors_copy = spqr@mydomain      postmaster@mydomain.example :\
                rpps@mydomain      hostmaster@mydomain.example, \
                                postmaster@mydomain.example

```

The address list is expanded before use. The expansion variables *\$local\_part* and *\$domain* are set from the original recipient of the error message, and if there was any wildcard matching in the pattern, the expansion variables *\$0*, *\$1*, etc. are set in the normal way.

<b>errors_reply_to</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
------------------------	------------------	---------------------	-----------------------

By default, Exim's bounce and delivery warning messages contain the header line

```
From: Mail Delivery System <Mailer-Daemon@qualify-domain>
```

where *qualify-domain* is the value of the **qualify\_domain** option. A warning message that is generated by the **quota\_warn\_message** option in an *appendfile* transport may contain its own *From:* header line that overrides the default.

Experience shows that people reply to bounce messages. If the **errors\_reply\_to** option is set, a *Reply-To:* header is added to bounce and warning messages. For example:

```
errors_reply_to = postmaster@my.domain.example
```

The value of the option is not expanded. It must specify a valid RFC 2822 address. However, if a warning message that is generated by the **quota\_warn\_message** option in an *appendfile* transport contain its own *Reply-To:* header line, the value of the **errors\_reply\_to** option is not used.

<b>event_action</b>	Use: <i>main</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
---------------------	------------------	----------------------------------	-----------------------

This option declares a string to be expanded for Exim's events mechanism. For details see chapter 60.

<b>exim_group</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>compile-time configured</i>
-------------------	------------------	---------------------	---

This option changes the gid under which Exim runs when it gives up root privilege. The default value is compiled into the binary. The value of this option is used only when **exim\_user** is also set. Unless it consists entirely of digits, the string is looked up using *getgrnam()*, and failure causes a configuration error. See chapter 55 for a discussion of security issues.

<b>exim_path</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>see below</i>
------------------	------------------	---------------------	---------------------------

This option specifies the path name of the Exim binary, which is used when Exim needs to re-exec itself. The default is set up to point to the file *exim* in the directory configured at compile time by the *BIN\_DIRECTORY* setting. It is necessary to change **exim\_path** if, exceptionally, Exim is run from some other place. **Warning:** Do not use a macro to define the value of this option, because you will break those Exim utilities that scan the configuration file to find where the binary is. (They then use the **-bP** option to extract option settings such as the value of **spool\_directory**.)

<b>exim_user</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>compile-time configured</i>
------------------	------------------	---------------------	---

This option changes the uid under which Exim runs when it gives up root privilege. The default value is compiled into the binary. Ownership of the run time configuration file and the use of the **-C** and **-D** command line options is checked against the values in the binary, not what is set here.

Unless it consists entirely of digits, the string is looked up using *getpwnam()*, and failure causes a configuration error. If **exim\_group** is not also supplied, the gid is taken from the result of *getpwnam()* if it is used. See chapter 55 for a discussion of security issues.

<b>extra_local_interfaces</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>unset</i>
-------------------------------	------------------	--------------------------	-----------------------

This option defines network interfaces that are to be considered local when routing, but which are not used for listening by the daemon. See section 13.8 for details.

<b>extract_addresses_remove_arguments</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
---	------------------	----------------------	----------------------

According to some Sendmail documentation (Sun, IRIX, HP-UX), if any addresses are present on the command line when the **-t** option is used to build an envelope from a message's *To:*, *Cc:* and *Bcc:* headers, the command line addresses are removed from the recipients list. This is also how Smail behaves. However, other Sendmail documentation (the O'Reilly book) states that command line addresses are added to those obtained from the header lines. When **extract\_addresses\_remove\_arguments** is true (the default), Exim subtracts argument headers. If it is set false, Exim adds rather than removes argument addresses.

<b>finduser_retries</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>0</i>
-------------------------	------------------	----------------------	-------------------

On systems running NIS or other schemes in which user and group information is distributed from a remote system, there can be times when *getpwnam()* and related functions fail, even when given valid data, because things time out. Unfortunately these failures cannot be distinguished from genuine “not found” errors. If **finduser\_retries** is set greater than zero, Exim will try that many extra times to find a user or a group, waiting for one second between retries.

You should not set this option greater than zero if your user information is in a traditional */etc/passwd* file, because it will cause Exim needlessly to search the file multiple times for non-existent users, and also cause delay.

<b>freeze_tell</b>	Use: <i>main</i>	Type: <i>string list, comma separated</i>	Default: <i>unset</i>
--------------------	------------------	---	-----------------------

On encountering certain errors, or when configured to do so in a system filter, ACL, or special router, Exim freezes a message. This means that no further delivery attempts take place until an administrator thaws the message, or the **auto\_thaw**, **ignore\_bounce\_errors\_after**, or **timeout\_frozen\_after** feature cause it to be processed. If **freeze\_tell** is set, Exim generates a warning message whenever it freezes something, unless the message it is freezing is a locally-generated bounce message. (Without this exception there is the possibility of looping.) The warning message is sent to the addresses supplied as the comma-separated value of this option. If several of the message's addresses cause freezing, only a single message is sent. If the freezing was automatic, the reason(s) for freezing can be found in the message log. If you configure freezing in a filter or ACL, you must arrange for any logging that you require.

<b>gecos_name</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------	------------------	----------------------	-----------------------

Some operating systems, notably HP-UX, use the “gecos” field in the system password file to hold other information in addition to users' real names. Exim looks up this field for use when it is creating *Sender:* or *From:* headers. If either **gecos\_pattern** or **gecos\_name** are unset, the contents of the field are used unchanged, except that, if an ampersand is encountered, it is replaced by the user's login name with the first character forced to upper case, since this is a convention that is observed on many systems.

When these options are set, **gecos\_pattern** is treated as a regular expression that is to be applied to the field (again with & replaced by the login name), and if it matches, **gecos\_name** is expanded and used as the user's name.

Numeric variables such as \$1, \$2, etc. can be used in the expansion to pick up sub-fields that were matched by the pattern. In HP-UX, where the user's name terminates at the first comma, the following can be used:

```
gecos_pattern = ([^,]*)
gecos_name = $1
```

<b>gecos_pattern</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------	------------------	---------------------	-----------------------

See **gecos\_name** above.

<b>gnutls_compat_mode</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>unset</i>
---------------------------	------------------	----------------------	-----------------------

This option controls whether GnuTLS is used in compatibility mode in an Exim server. This reduces security slightly, but improves interworking with older implementations of TLS.

option gnutls\_allow\_auto\_pkcs11 main boolean unset This option will let GnuTLS (2.12.0 or later) autoload PKCS11 modules with the p11-kit configuration files in */etc/pkcs11/modules/*.

See <http://www.gnutls.org/manual/gnutls.html#Smart-cards-and-HSMs> for documentation.

<b>headers_charset</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>see below</i>
------------------------	------------------	---------------------	---------------------------

This option sets a default character set for translating from encoded MIME “words” in header lines, when referenced by an *\$h\_xxx* expansion item. The default is the value of **HEADERS\_CHARSET** in *Local/Makefile*. The ultimate default is ISO-8859-1. For more details see the description of header insertions in section 11.5.

<b>header_maxsize</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>see below</i>
-----------------------	------------------	----------------------	---------------------------

This option controls the overall maximum size of a message's header section. The default is the value of **HEADER\_MAXSIZE** in *Local/Makefile*; the default for that is 1M. Messages with larger header sections are rejected.

<b>header_line_maxsize</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>0</i>
----------------------------	------------------	----------------------	-------------------

This option limits the length of any individual header line in a message, after all the continuations have been joined together. Messages with individual header lines that are longer than the limit are rejected. The default value of zero means “no limit”.

<b>helo_accept_junk_hosts</b>	Use: <i>main</i>	Type: <i>host list</i> <sup>†</sup>	Default: <i>unset</i>
-------------------------------	------------------	-------------------------------------	-----------------------

Exim checks the syntax of HELO and EHLO commands for incoming SMTP mail, and gives an error response for invalid data. Unfortunately, there are some SMTP clients that send syntactic junk. They can be accommodated by setting this option. Note that this is a syntax check only. See **helo\_verify\_hosts** if you want to do semantic checking. See also **helo\_allow\_chars** for a way of extending the permitted character set.

<b>helo_allow_chars</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------------	------------------	---------------------	-----------------------

This option can be set to a string of rogue characters that are permitted in all EHLO and HELO names in addition to the standard letters, digits, hyphens, and dots. If you really must allow underscores, you can set

```
helo_allow_chars = _
```

Note that the value is one string, not a list.

<b>helo_lookup_domains</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>@ : @ [ ]</i>
----------------------------	------------------	---------------------------	---------------------------

If the domain given by a client in a HELO or EHLO command matches this list, a reverse lookup is done in order to establish the host's true name. The default forces a lookup if the client host gives the server's name or any of its IP addresses (in brackets), something that broken clients have been seen to do.

<b>helo_try_verify_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>unset</i>
------------------------------	------------------	-------------------------	-----------------------

By default, Exim just checks the syntax of HELO and EHLO commands (see **helo\_accept\_junk\_hosts** and **helo\_allow\_chars**). However, some sites like to do more extensive checking of the data supplied by these commands. The ACL condition `verify = helo` is provided to make this possible. Formerly, it was necessary also to set this option (**helo\_try\_verify\_hosts**) to force the check to occur. From release 4.53 onwards, this is no longer necessary. If the check has not been done before `verify = helo` is encountered, it is done at that time. Consequently, this option is obsolete. Its specification is retained here for backwards compatibility.

When an EHLO or HELO command is received, if the calling host matches **helo\_try\_verify\_hosts**, Exim checks that the host name given in the HELO or EHLO command either:

- is an IP literal matching the calling address of the host, or
- matches the host name that Exim obtains by doing a reverse lookup of the calling host address, or
- when looked up in DNS yields the calling host address.

However, the EHLO or HELO command is not rejected if any of the checks fail. Processing continues, but the result of the check is remembered, and can be detected later in an ACL by the `verify = helo` condition.

If DNS was used for successful verification, the variable `$helo_verify_dnssec` records the DNSSEC status of the lookups.

<b>helo_verify_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>unset</i>
--------------------------	------------------	-------------------------	-----------------------

Like **helo\_try\_verify\_hosts**, this option is obsolete, and retained only for backwards compatibility. For hosts that match this option, Exim checks the host name given in the HELO or EHLO in the same way as for **helo\_try\_verify\_hosts**. If the check fails, the HELO or EHLO command is rejected with a 550 error, and entries are written to the main and reject logs. If a MAIL command is received before EHLO or HELO, it is rejected with a 503 error.

<b>hold_domains</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
---------------------	------------------	---------------------------	-----------------------

This option allows mail for particular domains to be held on the queue manually. The option is overridden if a message delivery is forced with the **-M**, **-qf**, **-Rf** or **-Sf** options, and also while testing or verifying addresses using **-bt** or **-bv**. Otherwise, if a domain matches an item in **hold\_domains**, no routing or delivery for that address is done, and it is deferred every time the message is looked at.

This option is intended as a temporary operational measure for delaying the delivery of mail while some problem is being sorted out, or some new configuration tested. If you just want to delay the processing of some domains until a queue run occurs, you should use **queue\_domains** or **queue\_smtplib\_domains**, not **hold\_domains**.

A setting of **hold\_domains** does not override Exim's code for removing messages from the queue if they have been there longer than the longest retry time in any retry rule. If you want to hold messages for longer than the normal retry times, insert a dummy retry rule with a long retry time.

<b>host_lookup</b>	Use: <i>main</i>	Type: <i>host list</i> †	Default: <i>unset</i>
--------------------	------------------	--------------------------	-----------------------

Exim does not look up the name of a calling host from its IP address unless it is required to compare against some host list, or the host matches **helo\_try\_verify\_hosts** or **helo\_verify\_hosts**, or the host matches this option (which normally contains IP addresses rather than host names). The default configuration file contains

```
host_lookup = *
```

which causes a lookup to happen for all hosts. If the expense of these lookups is felt to be too great, the setting can be changed or removed.

After a successful reverse lookup, Exim does a forward lookup on the name it has obtained, to verify that it yields the IP address that it started with. If this check fails, Exim behaves as if the name lookup failed.

After any kind of failure, the host name (in *\$sender\_host\_name*) remains unset, and *\$host\_lookup\_failed* is set to the string “1”. See also **dns\_again\_means\_nonexist**, **helo\_lookup\_domains**, and `verify = reverse_host_lookup` in ACLs.

<b>host_lookup_order</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>bydns:byaddr</i>
--------------------------	------------------	--------------------------	---------------------------------

This option specifies the order of different lookup methods when Exim is trying to find a host name from an IP address. The default is to do a DNS lookup first, and then to try a local lookup (using *gethostbyaddr()* or equivalent) if that fails. You can change the order of these lookups, or omit one entirely, if you want.

**Warning:** The “byaddr” method does not always yield aliases when there are multiple PTR records in the DNS and the IP address is not listed in */etc/hosts*. Different operating systems give different results in this case. That is why the default tries a DNS lookup first.

<b>host_reject_connection</b>	Use: <i>main</i>	Type: <i>host list</i> †	Default: <i>unset</i>
-------------------------------	------------------	--------------------------	-----------------------

If this option is set, incoming SMTP calls from the hosts listed are rejected as soon as the connection is made. This option is obsolete, and retained only for backward compatibility, because nowadays the ACL specified by **acl\_smtp\_connect** can also reject incoming connections immediately.

The ability to give an immediate rejection (either by this option or using an ACL) is provided for use in unusual cases. Many hosts will just try again, sometimes without much delay. Normally, it is better to use an ACL to reject incoming messages at a later stage, such as after RCPT commands. See chapter 43.

<b>hosts_connection_nolog</b>	Use: <i>main</i>	Type: <i>host list</i> †	Default: <i>unset</i>
-------------------------------	------------------	--------------------------	-----------------------

This option defines a list of hosts for which connection logging does not happen, even though the **smtp\_connection** log selector is set. For example, you might want not to log SMTP connections from local processes, or from 127.0.0.1, or from your local LAN. This option is consulted in the main loop of the daemon; you should therefore strive to restrict its value to a short inline list of IP addresses and networks. To disable logging SMTP connections from local processes, you must create a host list with an empty item. For example:

```
hosts_connection_nolog = :
```

If the **smtp\_connection** log selector is not set, this option has no effect.

<b>hosts_proxy</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>unset</i>
--------------------	------------------	-------------------------	-----------------------

This option enables use of Proxy Protocol proxies for incoming connections. For details see section 58.1.

<b>hosts_treat_as_local</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
-----------------------------	------------------	---------------------------	-----------------------

If this option is set, any host names that match the domain list are treated as if they were the local host when Exim is scanning host lists obtained from MX records or other sources. Note that the value of this option is a domain list, not a host list, because it is always used to check host names, not IP addresses.

This option also applies when Exim is matching the special items @mx\_any, @mx\_primary, and @mx\_secondary in a domain list (see section 10.8), and when checking the **hosts** option in the *smtp* transport for the local host (see the **allow\_localhost** option in that transport). See also **local\_interfaces**, **extra\_local\_interfaces**, and chapter 13, which contains a discussion about local network interfaces and recognizing the local host.

<b>ibase_servers</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>unset</i>
----------------------	------------------	--------------------------	-----------------------

This option provides a list of InterBase servers and associated connection data, to be used in conjunction with *ibase* lookups (see section 9.22). The option is available only if Exim has been built with InterBase support.

<b>ignore_bounce_errors_after</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>10w</i>
-----------------------------------	------------------	-------------------	---------------------

This option affects the processing of bounce messages that cannot be delivered, that is, those that suffer a permanent delivery failure. (Bounce messages that suffer temporary delivery failures are of course retried in the usual way.)

After a permanent delivery failure, bounce messages are frozen, because there is no sender to whom they can be returned. When a frozen bounce message has been on the queue for more than the given time, it is unfrozen at the next queue run, and a further delivery is attempted. If delivery fails again, the bounce message is discarded. This makes it possible to keep failed bounce messages around for a shorter time than the normal maximum retry time for frozen messages. For example,

```
ignore_bounce_errors_after = 12h
```

retries failed bounce message deliveries after 12 hours, discarding any further failures. If the value of this option is set to a zero time period, bounce failures are discarded immediately. Setting a very long time (as in the default value) has the effect of disabling this option. For ways of automatically dealing with other kinds of frozen message, see **auto\_thaw** and **timeout\_frozen\_after**.

<b>ignore_fromline_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>unset</i>
------------------------------	------------------	-------------------------	-----------------------

Some broken SMTP clients insist on sending a UUCP-like “From ” line before the headers of a message. By default this is treated as the start of the message’s body, which means that any following headers are not recognized as such. Exim can be made to ignore it by setting **ignore\_fromline\_hosts** to match those hosts that insist on sending it. If the sender is actually a local process rather than a remote host, and is using **-bs** to inject the messages, **ignore\_fromline\_local** must be set to achieve this effect.

<b>ignore_fromline_local</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------------	------------------	----------------------	-----------------------

See **ignore\_fromline\_hosts** above.

<b>keep_environment</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>unset</i>
-------------------------	------------------	--------------------------	-----------------------

This option contains a string list of environment variables to keep. You have to trust these variables or you have to be sure that these variables do not impose any security risk. Keep in mind that during the startup phase Exim is running with an effective UID 0 in most installations. As the default value is an empty list, the default environment for using libraries, running embedded Perl code, or running external binaries is empty, and does not even contain PATH or HOME.

Actually the list is interpreted as a list of patterns (10.1), except that it is not expanded first.

WARNING: Macro substitution is still done first, so having a macro FOO and having FOO\_HOME in your **keep\_environment** option may have unexpected results. You may work around this using a regular expression that does not match the macro name: `^[F]OO_HOME$`.

Current versions of Exim issue a warning during startup if you do not mention **keep\_environment** in your runtime configuration file and if your current environment is not empty. Future versions may not issue that warning anymore.

See the **add\_environment** main config option for a way to set environment variables to a fixed value. The environment for *pipe* transports is handled separately, see section 29.4 for details.

<b>keep_malformed</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>4d</i>
-----------------------	------------------	-------------------	--------------------

This option specifies the length of time to keep messages whose spool files have been corrupted in some way. This should, of course, never happen. At the next attempt to deliver such a message, it gets removed. The incident is logged.

<b>ldap_ca_cert_dir</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------------	------------------	---------------------	-----------------------

This option indicates which directory contains CA certificates for verifying a TLS certificate presented by an LDAP server. While Exim does not provide a default value, your SSL library may. Analogous to **tls\_verify\_certificates** but as a client-side option for LDAP and constrained to be a directory.

<b>ldap_ca_cert_file</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------------------	------------------	---------------------	-----------------------

This option indicates which file contains CA certificates for verifying a TLS certificate presented by an LDAP server. While Exim does not provide a default value, your SSL library may. Analogous to **tls\_verify\_certificates** but as a client-side option for LDAP and constrained to be a file.

<b>ldap_cert_file</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
-----------------------	------------------	---------------------	-----------------------

This option indicates which file contains an TLS client certificate which Exim should present to the LDAP server during TLS negotiation. Should be used together with **ldap\_cert\_key**.

<b>ldap_cert_key</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------	------------------	---------------------	-----------------------

This option indicates which file contains the secret/private key to use to prove identity to the LDAP server during TLS negotiation. Should be used together with **ldap\_cert\_file**, which contains the identity to be proven.

<b>ldap_cipher_suite</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------------------	------------------	---------------------	-----------------------

This controls the TLS cipher-suite negotiation during TLS negotiation with the LDAP server. See 42.4 for more details of the format of cipher-suite options with OpenSSL (as used by LDAP client libraries).

<b>ldap_default_servers</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>unset</i>
-----------------------------	------------------	--------------------------	-----------------------

This option provides a list of LDAP servers which are tried in turn when an LDAP query does not contain a server. See section 9.15 for details of LDAP queries. This option is available only when Exim has been built with LDAP support.

<b>ldap_require_cert</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i> .
--------------------------	------------------	---------------------	-------------------------

This should be one of the values "hard", "demand", "allow", "try" or "never". A value other than one of these is interpreted as "never". See the entry "TLS\_REQCERT" in your system man page for ldap.conf(5). Although Exim does not set a default, the LDAP library probably defaults to hard/demand.

<b>ldap_start_tls</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	------------------	----------------------	-----------------------

If set, Exim will attempt to negotiate TLS with the LDAP server when connecting on a regular LDAP port. This is the LDAP equivalent of SMTP's "STARTTLS". This is distinct from using "ldaps", which is the LDAP form of SSL-on-connect. In the event of failure to negotiate TLS, the action taken is controlled by **ldap\_require\_cert**. This option is ignored for `ldapi` connections.

<b>ldap_version</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>unset</i>
---------------------	------------------	----------------------	-----------------------

This option can be used to force Exim to set a specific protocol version for LDAP. If it option is unset, it is shown by the **-bP** command line option as -1. When this is the case, the default is 3 if LDAP\_VERSION3 is defined in the LDAP headers; otherwise it is 2. This option is available only when Exim has been built with LDAP support.

<b>local_from_check</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
-------------------------	------------------	----------------------	----------------------

When a message is submitted locally (that is, not over a TCP/IP connection) by an untrusted user, Exim removes any existing *Sender:* header line, and checks that the *From:* header line matches the login of the calling user and the domain specified by **qualify\_domain**.

**Note:** An unqualified address (no domain) in the *From:* header in a locally submitted message is automatically qualified by Exim, unless the **-bnq** command line option is used.

You can use **local\_from\_prefix** and **local\_from\_suffix** to permit affixes on the local part. If the *From:* header line does not match, Exim adds a *Sender:* header with an address constructed from the calling user's login and the default qualify domain.

If **local\_from\_check** is set false, the *From:* header check is disabled, and no *Sender:* header is ever added. If, in addition, you want to retain *Sender:* header lines supplied by untrusted users, you must also set **local\_sender\_retain** to be true.

These options affect only the header lines in the message. The envelope sender is still forced to be the login id at the qualify domain unless **untrusted\_set\_sender** permits the user to supply an envelope sender.

For messages received over TCP/IP, an ACL can specify "submission mode" to request similar header line checking. See section 47.16, which has more details about *Sender:* processing.

<b>local_from_prefix</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------------------	------------------	---------------------	-----------------------

When Exim checks the *From:* header line of locally submitted messages for matching the login id (see **local\_from\_check** above), it can be configured to ignore certain prefixes and suffixes in the local part of the address. This is done by setting **local\_from\_prefix** and/or **local\_from\_suffix** to appropri-



ate lists, in the same form as the **local\_part\_prefix** and **local\_part\_suffix** router options (see chapter 15). For example, if

```
local_from_prefix = *-
```

is set, a *From:* line containing

```
From: anything-user@your.domain.example
```

will not cause a *Sender:* header to be added if *user@your.domain.example* matches the actual sender address that is constructed from the login name and qualify domain.

<b>local_from_suffix</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------------------	------------------	---------------------	-----------------------

See **local\_from\_prefix** above.

<b>local_interfaces</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>see below</i>
-------------------------	------------------	--------------------------	---------------------------

This option controls which network interfaces are used by the daemon for listening; they are also used to identify the local host when routing. Chapter 13 contains a full description of this option and the related options **daemon\_smtp\_ports**, **extra\_local\_interfaces**, **hosts\_treat\_as\_local**, and **tls\_on\_connect\_ports**. The default value for **local\_interfaces** is

```
local_interfaces = 0.0.0.0
```

when Exim is built without IPv6 support; otherwise it is

```
local_interfaces = <; ::0 ; 0.0.0.0
```

<b>local_scan_timeout</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>5m</i>
---------------------------	------------------	-------------------	--------------------

This timeout applies to the *local\_scan()* function (see chapter 45). Zero means “no timeout”. If the timeout is exceeded, the incoming message is rejected with a temporary error if it is an SMTP message. For a non-SMTP message, the message is dropped and Exim ends with a non-zero code. The incident is logged on the main and reject logs.

<b>local_sender_retain</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------------	------------------	----------------------	-----------------------

When a message is submitted locally (that is, not over a TCP/IP connection) by an untrusted user, Exim removes any existing *Sender:* header line. If you do not want this to happen, you must set **local\_sender\_retain**, and you must also set **local\_from\_check** to be false (Exim will complain if you do not). See also the ACL modifier `control = suppress_local_fixups`. Section 47.16 has more details about *Sender:* processing.

<b>localhost_number</b>	Use: <i>main</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-------------------------	------------------	----------------------------------	-----------------------

Exim’s message ids are normally unique only within the local host. If uniqueness among a set of hosts is required, each host must set a different value for the **localhost\_number** option. The string is expanded immediately after reading the configuration file (so that a number can be computed from the host name, for example) and the result of the expansion must be a number in the range 0–16 (or 0–10 on operating systems with case-insensitive file systems). This is available in subsequent string expansions via the variable *\$localhost\_number*. When **localhost\_number** is set, the final two characters of the message id, instead of just being a fractional part of the time, are computed from the time and the local host number as described in section 3.4.

<b>log_file_path</b>	Use: <i>main</i>	Type: <i>string list†</i>	Default: <i>set at compile time</i>
----------------------	------------------	---------------------------	-------------------------------------

This option sets the path which is used to determine the names of Exim's log files, or indicates that logging is to be to syslog, or both. It is expanded when Exim is entered, so it can, for example, contain a reference to the host name. If no specific path is set for the log files at compile or run time, or if the option is unset at run time (i.e. `log_file_path =` ) they are written in a sub-directory called *log* in Exim's spool directory. Chapter 52 contains further details about Exim's logging, and section 52.1 describes how the contents of **log\_file\_path** are used. If this string is fixed at your installation (contains no expansion variables) it is recommended that you do not set this option in the configuration file, but instead supply the path using `LOG_FILE_PATH` in *Local/Makefile* so that it is available to Exim for logging errors detected early on – in particular, failure to read the configuration file.

<b>log_selector</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
---------------------	------------------	---------------------	-----------------------

This option can be used to reduce or increase the number of things that Exim writes to its log files. Its argument is made up of names preceded by plus or minus characters. For example:

```
log_selector = +arguments -retry_defer
```

A list of possible names and what they control is given in the chapter on logging, in section 52.15.

<b>log_timezone</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------	------------------	----------------------	-----------------------

By default, the timestamps on log lines are in local time without the timezone. This means that if your timezone changes twice a year, the timestamps in log lines are ambiguous for an hour when the clocks go back. One way of avoiding this problem is to set the timezone to UTC. An alternative is to set **log\_timezone** true. This turns on the addition of the timezone offset to timestamps in log lines. Turning on this option can add quite a lot to the size of log files because each line is extended by 6 characters. Note that the *\$tod\_log* variable contains the log timestamp without the zone, but there is another variable called *\$tod\_zone* that contains just the timezone offset.

<b>lookup_open_max</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: 25
------------------------	------------------	----------------------	-------------

This option limits the number of simultaneously open files for single-key lookups that use regular files (that is, *lsearch*, *dbm*, and *cdb*). Exim normally keeps these files open during routing, because often the same file is required several times. If the limit is reached, Exim closes the least recently used file. Note that if you are using the *ndbm* library, it actually opens two files for each logical DBM database, though it still counts as one for the purposes of **lookup\_open\_max**. If you are getting “too many open files” errors with NDBM, you need to reduce the value of **lookup\_open\_max**.

<b>max_username_length</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: 0
----------------------------	------------------	----------------------	------------

Some operating systems are broken in that they truncate long arguments to *getpwnam()* to eight characters, instead of returning “no such user”. If this option is set greater than zero, any attempt to call *getpwnam()* with an argument that is longer behaves as if *getpwnam()* failed.

<b>message_body_newlines</b>	Use: <i>main</i>	Type: <i>bool</i>	Default: <i>false</i>
------------------------------	------------------	-------------------	-----------------------

By default, newlines in the message body are replaced by spaces when setting the *\$message\_body* and *\$message\_body\_end* expansion variables. If this option is set true, this no longer happens.

<b>message_body_visible</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>500</i>
-----------------------------	------------------	----------------------	---------------------

This option specifies how much of a message's body is to be included in the *\$message\_body* and *\$message\_body\_end* expansion variables.

<b>message_id_header_domain</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
---------------------------------	------------------	----------------------	-----------------------

If this option is set, the string is expanded and used as the right hand side (domain) of the *Message-ID:* header that Exim creates if a locally-originated incoming message does not have one. "Locally-originated" means "not received over TCP/IP." Otherwise, the primary host name is used. Only letters, digits, dot and hyphen are accepted; any other characters are replaced by hyphens. If the expansion is forced to fail, or if the result is an empty string, the option is ignored.

<b>message_id_header_text</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------------	------------------	----------------------	-----------------------

If this variable is set, the string is expanded and used to augment the text of the *Message-id:* header that Exim creates if a locally-originated incoming message does not have one. The text of this header is required by RFC 2822 to take the form of an address. By default, Exim uses its internal message id as the local part, and the primary host name as the domain. If this option is set, it is expanded, and provided the expansion is not forced to fail, and does not yield an empty string, the result is inserted into the header immediately before the @, separated from the internal message id by a dot. Any characters that are illegal in an address are automatically converted into hyphens. This means that variables such as *\$tod\_log* can be used, because the spaces and colons will become hyphens.

<b>message_logs</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------------	------------------	----------------------	----------------------

If this option is turned off, per-message log files are not created in the *msglog* spool sub-directory. This reduces the amount of disk I/O required by Exim, by reducing the number of files involved in handling a message from a minimum of four (header spool file, body spool file, delivery journal, and per-message log) to three. The other major I/O activity is Exim's main log, which is not affected by this option.

<b>message_size_limit</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>50M</i>
---------------------------	------------------	----------------------	---------------------

This option limits the maximum size of message that Exim will process. The value is expanded for each incoming connection so, for example, it can be made to depend on the IP address of the remote host for messages arriving via TCP/IP. After expansion, the value must be a sequence of decimal digits, optionally followed by K or M.

**Note:** This limit cannot be made to depend on a message's sender or any other properties of an individual message, because it has to be advertised in the server's response to EHLO. String expansion failure causes a temporary error. A value of zero means no limit, but its use is not recommended. See also **bounce\_return\_size\_limit**.

Incoming SMTP messages are failed with a 552 error if the limit is exceeded; locally-generated messages either get a stderr message or a delivery failure message to the sender, depending on the **-oe** setting. Rejection of an oversized message is logged in both the main and the reject logs. See also the generic transport option **message\_size\_limit**, which limits the size of message that an individual transport can process.

If you use a virus-scanner and set this option to to a value larger than the maximum size that your virus-scanner is configured to support, you may get failures triggered by large mails. The right size to configure for the virus-scanner depends upon what data is passed and the options in use but it's probably safest to just set it to a little larger than this value. E.g., with a default Exim message size of 50M and a default ClamAV StreamMaxLength of 10M, some problems may result.

A value of 0 will disable size limit checking; Exim will still advertise the SIZE extension in an EHLO response, but without a limit, so as to permit SMTP clients to still indicate the message size along with the MAIL verb.

<b>move_frozen_messages</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------------	------------------	----------------------	-----------------------

This option, which is available only if Exim has been built with the setting

```
SUPPORT_MOVE_FROZEN_MESSAGES=yes
```

in *Local/Makefile*, causes frozen messages and their message logs to be moved from the *input* and *msglog* directories on the spool to *Finput* and *Fmsglog*, respectively. There is currently no support in Exim or the standard utilities for handling such moved messages, and they do not show up in lists generated by **-bp** or by the Exim monitor.

<b>mua_wrapper</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	------------------	----------------------	-----------------------

Setting this option true causes Exim to run in a very restrictive mode in which it passes messages synchronously to a smart host. Chapter 51 contains a full description of this facility.

<b>mysql_servers</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>unset</i>
----------------------	------------------	--------------------------	-----------------------

This option provides a list of MySQL servers and associated connection data, to be used in conjunction with *mysql* lookups (see section 9.22). The option is available only if Exim has been built with MySQL support.

<b>never_users</b>	Use: <i>main</i>	Type: <i>string list</i> <sup>†</sup>	Default: <i>unset</i>
--------------------	------------------	---------------------------------------	-----------------------

This option is expanded just once, at the start of Exim's processing. Local message deliveries are normally run in processes that are setuid to the recipient, and remote deliveries are normally run under Exim's own uid and gid. It is usually desirable to prevent any deliveries from running as root, as a safety precaution.

When Exim is built, an option called FIXED\_NEVER\_USERS can be set to a list of users that must not be used for local deliveries. This list is fixed in the binary and cannot be overridden by the configuration file. By default, it contains just the single user name "root". The **never\_users** runtime option can be used to add more users to the fixed list.

If a message is to be delivered as one of the users on the fixed list or the **never\_users** list, an error occurs, and delivery is deferred. A common example is

```
never_users = root:daemon:bin
```

Including root is redundant if it is also on the fixed list, but it does no harm. This option overrides the **pipe\_as\_creator** option of the *pipe* transport driver.

<b>openssl_options</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>+no_sslv2</i> <i>+single_dh_use</i> <i>+no_ticket</i>
------------------------	------------------	--------------------------	---

This option allows an administrator to adjust the SSL options applied by OpenSSL to connections. It is given as a space-separated list of items, each one to be +added or -subtracted from the current value.

This option is only available if Exim is built against OpenSSL. The values available for this option vary according to the age of your OpenSSL install. The "all" value controls a subset of flags which are available, typically the bug workaround options. The *SSL\_CTX\_set\_options* man page will list the

values known on your system and Exim should support all the “bug workaround” options and many of the “modifying” options. The Exim names lose the leading “SSL\_OP\_” and are lower-cased.

Note that adjusting the options can have severe impact upon the security of SSL as used by Exim. It is possible to disable safety checks and shoot yourself in the foot in various unpleasant ways. This option should not be adjusted lightly. An unrecognised item will be detected at startup, by invoking Exim with the **-bV** flag.

The option affects Exim operating both as a server and as a client.

Historical note: prior to release 4.80, Exim defaulted this value to “+dont\_insert\_empty\_fragments”, which may still be needed for compatibility with some clients, but which lowers security by increasing exposure to some now infamous attacks.

Examples:

```
# Make both old MS and old Eudora happy:
openssl_options = -all +microsoft_big_sslv3_buffer \
                  +dont_insert_empty_fragments

# Disable older protocol versions:
openssl_options = +no_sslv2 +no_sslv3
```

Possible options may include:

- all
- allow\_unsafe\_legacy\_renegotiation
- cipher\_server\_preference
- dont\_insert\_empty\_fragments
- ephemeral\_rsa
- legacy\_server\_connect
- microsoft\_big\_sslv3\_buffer
- microsoft\_sess\_id\_bug
- msie\_sslv2\_rsa\_padding
- netscape\_challenge\_bug
- netscape\_reuse\_cipher\_change\_bug
- no\_compression
- no\_session\_resumption\_on\_renegotiation
- no\_sslv2
- no\_sslv3
- no\_ticket
- no\_tlsv1
- no\_tlsv1\_1
- no\_tlsv1\_2
- safari\_ecdhe\_ecdsa\_bug
- single\_dh\_use
- single\_ecdh\_use
- ssleay\_080\_client\_dh\_bug
- sslref2\_reuse\_cert\_type\_bug

- `tls_block_padding_bug`
- `tls_d5_bug`
- `tls_rollback_bug`

As an aside, the `safari_ecdhe_ecdsa_bug` item is a misnomer and affects all clients connecting using the MacOS SecureTransport TLS facility prior to MacOS 10.8.4, including email clients. If you see old MacOS clients failing to negotiate TLS then this option value might help, provided that your OpenSSL release is new enough to contain this work-around. This may be a situation where you have to upgrade OpenSSL to get buggy clients working.

<b>oracle_servers</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>unset</i>
-----------------------	------------------	--------------------------	-----------------------

This option provides a list of Oracle servers and associated connection data, to be used in conjunction with *oracle* lookups (see section 9.22). The option is available only if Exim has been built with Oracle support.

<b>percent_hack_domains</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
-----------------------------	------------------	---------------------------	-----------------------

The “percent hack” is the convention whereby a local part containing a percent sign is re-interpreted as a new email address, with the percent replaced by @. This is sometimes called “source routing”, though that term is also applied to RFC 2822 addresses that begin with an @ character. If this option is set, Exim implements the percent facility for those domains listed, but no others. This happens before an incoming SMTP address is tested against an ACL.

**Warning:** The “percent hack” has often been abused by people who are trying to get round relaying restrictions. For this reason, it is best avoided if at all possible. Unfortunately, a number of less security-conscious MTAs implement it unconditionally. If you are running Exim on a gateway host, and routing mail through to internal MTAs without processing the local parts, it is a good idea to reject recipient addresses with percent characters in their local parts. Exim’s default configuration does this.

<b>perl_at_start</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------	----------------------	-----------------------

This option is available only when Exim is built with an embedded Perl interpreter. See chapter 12 for details of its use.

<b>perl_startup</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
---------------------	------------------	---------------------	-----------------------

This option is available only when Exim is built with an embedded Perl interpreter. See chapter 12 for details of its use.

<b>perl_startup</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------	------------------	----------------------	-----------------------

This Option enables the taint mode of the embedded Perl interpreter.

<b>pgsql_servers</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>unset</i>
----------------------	------------------	--------------------------	-----------------------

This option provides a list of PostgreSQL servers and associated connection data, to be used in conjunction with *pgsql* lookups (see section 9.22). The option is available only if Exim has been built with PostgreSQL support.

<b>pid_file_path</b>	Use: <i>main</i>	Type: <i>string</i> †	Default: <i>set at compile time</i>
----------------------	------------------	-----------------------	-------------------------------------

This option sets the name of the file to which the Exim daemon writes its process id. The string is expanded, so it can contain, for example, references to the host name:

```
pid_file_path = /var/log/$primary_hostname/exim.pid
```

If no path is set, the pid is written to the file *exim-daemon.pid* in Exim’s spool directory. The value set by the option can be overridden by the **-oP** command line option. A pid file is not written if a “non-standard” daemon is run by means of the **-oX** option, unless a path is explicitly supplied by **-oP**.

<b>pipelining_advertise_hosts</b>	Use: <i>main</i>	Type: <i>host list</i> †	Default: <i>*</i>
-----------------------------------	------------------	--------------------------	-------------------

This option can be used to suppress the advertisement of the SMTP PIPELINING extension to specific hosts. See also the **no\_pipelining** control in section 43.22. When PIPELINING is not advertised and **smtp\_enforce\_sync** is true, an Exim server enforces strict synchronization for each SMTP command and response. When PIPELINING is advertised, Exim assumes that clients will use it; “out of order” commands that are “expected” do not count as protocol errors (see **smtp\_max\_synprot\_errors**).

<b>prdr_enable</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	------------------	----------------------	-----------------------

This option can be used to enable the Per-Recipient Data Response extension to SMTP, defined by Eric Hall. If the option is set, PRDR is advertised by Exim when operating as a server. If the client requests PRDR, and more than one recipient, for a message an additional ACL is called for each recipient after the message content is received. See section 43.9.

<b>preserve_message_logs</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------------	------------------	----------------------	-----------------------

If this option is set, message log files are not deleted when messages are completed. Instead, they are moved to a sub-directory of the spool directory called *msglog.OLD*, where they remain available for statistical or debugging purposes. This is a dangerous option to set on systems with any appreciable volume of mail. Use with care!

<b>primary_hostname</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>see below</i>
-------------------------	------------------	---------------------	---------------------------

This specifies the name of the current host. It is used in the default EHLO or HELO command for outgoing SMTP messages (changeable via the **helo\_data** option in the *smtp* transport), and as the default for **qualify\_domain**. The value is also used by default in some SMTP response messages from an Exim server. This can be changed dynamically by setting **smtp\_active\_hostname**.

If **primary\_hostname** is not set, Exim calls *uname()* to find the host name. If this fails, Exim panics and dies. If the name returned by *uname()* contains only one component, Exim passes it to *gethostbyname()* (or *getipnodebyname()* when available) in order to obtain the fully qualified version. The variable *\$primary\_hostname* contains the host name, whether set explicitly by this option, or defaulted.

<b>print_topbitchars</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------------	------------------	----------------------	-----------------------

By default, Exim considers only those characters whose codes lie in the range 32–126 to be printing characters. In a number of circumstances (for example, when writing log entries) non-printing characters are converted into escape sequences, primarily to avoid messing up the layout. If **print\_topbitchars** is set, code values of 128 and above are also considered to be printing characters.

This option also affects the header syntax checks performed by the *autoreply* transport, and whether Exim uses RFC 2047 encoding of the user's full name when constructing From: and Sender: addresses (as described in section 47.18). Setting this option can cause Exim to generate eight bit message headers that do not conform to the standards.

<b>process_log_path</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------------	------------------	---------------------	-----------------------

This option sets the name of the file to which an Exim process writes its “process log” when sent a USR1 signal. This is used by the *exiwhat* utility script. If this option is unset, the file called *exim-process.info* in Exim's spool directory is used. The ability to specify the name explicitly can be useful in environments where two different Exims are running, using different spool directories.

<b>prod_requires_admin</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
----------------------------	------------------	----------------------	----------------------

The **-M**, **-R**, and **-q** command-line options require the caller to be an admin user unless **prod\_requires\_admin** is set false. See also **queue\_list\_requires\_admin** and **commandline\_checks\_require\_admin**.

<b>qualify_domain</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>see below</i>
-----------------------	------------------	---------------------	---------------------------

This option specifies the domain name that is added to any envelope sender addresses that do not have a domain qualification. It also applies to recipient addresses if **qualify\_recipient** is not set. Unqualified addresses are accepted by default only for locally-generated messages. Qualification is also applied to addresses in header lines such as *From:* and *To:* for locally-generated messages, unless the **-bnq** command line option is used.

Messages from external sources must always contain fully qualified addresses, unless the sending host matches **sender\_unqualified\_hosts** or **recipient\_unqualified\_hosts** (as appropriate), in which case incoming addresses are qualified with **qualify\_domain** or **qualify\_recipient** as necessary. Internally, Exim always works with fully qualified envelope addresses. If **qualify\_domain** is not set, it defaults to the **primary\_hostname** value.

<b>qualify_recipient</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>see below</i>
--------------------------	------------------	---------------------	---------------------------

This option allows you to specify a different domain for qualifying recipient addresses to the one that is used for senders. See **qualify\_domain** above.

<b>queue_domains</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
----------------------	------------------	---------------------------	-----------------------

This option lists domains for which immediate delivery is not required. A delivery process is started whenever a message is received, but only those domains that do not match are processed. All other deliveries wait until the next queue run. See also **hold\_domains** and **queue\_smtp\_domains**.

<b>queue_list_requires_admin</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
----------------------------------	------------------	----------------------	----------------------

The **-bp** command-line option, which lists the messages that are on the queue, requires the caller to be an admin user unless **queue\_list\_requires\_admin** is set false. See also **prod\_requires\_admin** and **commandline\_checks\_require\_admin**.

<b>queue_only</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------	------------------	----------------------	-----------------------

If **queue\_only** is set, a delivery process is not automatically started whenever a message is received. Instead, the message waits on the queue for the next queue run. Even if **queue\_only** is false, incoming messages may not get delivered immediately when certain conditions (such as heavy load) occur.



The **-odq** command line has the same effect as **queue\_only**. The **-odb** and **-odi** command line options override **queue\_only** unless **queue\_only\_override** is set false. See also **queue\_only\_file**, **queue\_only\_load**, and **smtp\_accept\_queue**.

<b>queue_only_file</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
------------------------	------------------	---------------------	-----------------------

This option can be set to a colon-separated list of absolute path names, each one optionally preceded by “smtp”. When Exim is receiving a message, it tests for the existence of each listed path using a call to *stat()*. For each path that exists, the corresponding queueing option is set. For paths with no prefix, **queue\_only** is set; for paths prefixed by “smtp”, **queue\_smtp\_domains** is set to match all domains. So, for example,

```
queue_only_file = smtp/some/file
```

causes Exim to behave as if **queue\_smtp\_domains** were set to “\*” whenever */some/file* exists.

<b>queue_only_load</b>	Use: <i>main</i>	Type: <i>fixed-point</i>	Default: <i>unset</i>
------------------------	------------------	--------------------------	-----------------------

If the system load average is higher than this value, incoming messages from all sources are queued, and no automatic deliveries are started. If this happens during local or remote SMTP input, all subsequent messages received on the same SMTP connection are queued by default, whatever happens to the load in the meantime, but this can be changed by setting **queue\_only\_load\_latch** false.

Deliveries will subsequently be performed by queue runner processes. This option has no effect on ancient operating systems on which Exim cannot determine the load average. See also **deliver\_queue\_load\_max** and **smtp\_load\_reserve**.

<b>queue_only_load_latch</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
------------------------------	------------------	----------------------	----------------------

When this option is true (the default), once one message has been queued because the load average is higher than the value set by **queue\_only\_load**, all subsequent messages received on the same SMTP connection are also queued. This is a deliberate choice; even though the load average may fall below the threshold, it doesn’t seem right to deliver later messages on the same connection when not delivering earlier ones. However, there are special circumstances such as very long-lived connections from scanning appliances where this is not the best strategy. In such cases, **queue\_only\_load\_latch** should be set false. This causes the value of the load average to be re-evaluated for each message.

<b>queue_only_override</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
----------------------------	------------------	----------------------	----------------------

When this option is true, the **-odx** command line options override the setting of **queue\_only** or **queue\_only\_file** in the configuration file. If **queue\_only\_override** is set false, the **-odx** options cannot be used to override; they are accepted, but ignored.

<b>queue_run_in_order</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	------------------	----------------------	-----------------------

If this option is set, queue runs happen in order of message arrival instead of in an arbitrary order. For this to happen, a complete list of the entire queue must be set up before the deliveries start. When the queue is all held in a single directory (the default), a single list is created for both the ordered and the non-ordered cases. However, if **split\_spool\_directory** is set, a single list is not created when **queue\_run\_in\_order** is false. In this case, the sub-directories are processed one at a time (in a random order), and this avoids setting up one huge list for the whole queue. Thus, setting **queue\_run\_in\_order** with **split\_spool\_directory** may degrade performance when the queue is large, because of the extra work in setting up the single, large list. In most situations, **queue\_run\_in\_order** should not be set.

<b>queue_run_max</b>	Use: <i>main</i>	Type: <i>integer</i> †	Default: 5
----------------------	------------------	------------------------	------------

This controls the maximum number of queue runner processes that an Exim daemon can run simultaneously. This does not mean that it starts them all at once, but rather that if the maximum number are still running when the time comes to start another one, it refrains from starting another one. This can happen with very large queues and/or very sluggish deliveries. This option does not, however, interlock with other processes, so additional queue runners can be started by other means, or by killing and restarting the daemon.

Setting this option to zero does not suppress queue runs; rather, it disables the limit, allowing any number of simultaneous queue runner processes to be run. If you do not want queue runs to occur, omit the **-qxx** setting on the daemon's command line.

To set limits for different named queues use an expansion depending on the *\$queue\_name* variable.

<b>queue_smtp_domains</b>	Use: <i>main</i>	Type: <i>domain list</i> †	Default: <i>unset</i>
---------------------------	------------------	----------------------------	-----------------------

When this option is set, a delivery process is started whenever a message is received, routing is performed, and local deliveries take place. However, if any SMTP deliveries are required for domains that match **queue\_smtp\_domains**, they are not immediately delivered, but instead the message waits on the queue for the next queue run. Since routing of the message has taken place, Exim knows to which remote hosts it must be delivered, and so when the queue run happens, multiple messages for the same host are delivered over a single SMTP connection. The **-odqs** command line option causes all SMTP deliveries to be queued in this way, and is equivalent to setting **queue\_smtp\_domains** to **"\*"**. See also **hold\_domains** and **queue\_domains**.

<b>receive_timeout</b>	Use: <i>main</i>	Type: <i>time</i>	Default: 0s
------------------------	------------------	-------------------	-------------

This option sets the timeout for accepting a non-SMTP message, that is, the maximum time that Exim waits when reading a message on the standard input. If the value is zero, it will wait for ever. This setting is overridden by the **-or** command line option. The timeout for incoming SMTP messages is controlled by **smtp\_receive\_timeout**.

<b>received_header_text</b>	Use: <i>main</i>	Type: <i>string</i> †	Default: <i>see below</i>
-----------------------------	------------------	-----------------------	---------------------------

This string defines the contents of the *Received:* message header that is added to each message, except for the timestamp, which is automatically added on at the end (preceded by a semicolon). The string is expanded each time it is used. If the expansion yields an empty string, no *Received:* header line is added to the message. Otherwise, the string should start with the text "Received:" and conform to the RFC 2822 specification for *Received:* header lines. The default setting is:

```
received_header_text = Received: \
  ${if def:sender_rcvhost {from $sender_rcvhost\n\t}\
  ${if def:sender_ident \
  {from ${quote_local_part:$sender_ident} }}\
  ${if def:sender_helo_name {(helo=$sender_helo_name)\n\t}}}\
  by $primary_hostname \
  ${if def:received_protocol {with $received_protocol}} \
  ${if def:tls_in_cipher {($tls_in_cipher)\n\t}}\
  (Exim $version_number)\n\t\
  ${if def:sender_address \
  {(envelope-from <$sender_address>)\n\t}}\
  id $message_exim_id\
  ${if def:received_for {\n\tfor $received_for}}
```

The reference to the TLS cipher is omitted when Exim is built without TLS support. The use of conditional expansions ensures that this works for both locally generated messages and messages received from remote hosts, giving header lines such as the following:

```
Received: from scrooge.carol.example ([192.168.12.25] ident=root)
by marley.carol.example with esmtp (Exim 4.00)
(envelope-from <bob@carol.example>)
id 16IOWa-000191-00
for chas@dickens.example; Tue, 25 Dec 2001 14:43:44 +0000
Received: by scrooge.carol.example with local (Exim 4.00)
id 16IOWW-000083-00; Tue, 25 Dec 2001 14:43:41 +0000
```

Until the body of the message has been received, the timestamp is the time when the message started to be received. Once the body has arrived, and all policy checks have taken place, the timestamp is updated to the time at which the message was accepted.

<b>received_headers_max</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>30</i>
-----------------------------	------------------	----------------------	--------------------

When a message is to be delivered, the number of *Received:* headers is counted, and if it is greater than this parameter, a mail loop is assumed to have occurred, the delivery is abandoned, and an error message is generated. This applies to both local and remote deliveries.

<b>recipient_unqualified_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>unset</i>
------------------------------------	------------------	-------------------------	-----------------------

This option lists those hosts from which Exim is prepared to accept unqualified recipient addresses in message envelopes. The addresses are made fully qualified by the addition of the **qualify\_recipient** value. This option also affects message header lines. Exim does not reject unqualified recipient addresses in headers, but it qualifies them only if the message came from a host that matches **recipient\_unqualified\_hosts**, or if the message was submitted locally (not using TCP/IP), and the **-bnq** option was not set.

<b>recipients_max</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>0</i>
-----------------------	------------------	----------------------	-------------------

If this option is set greater than zero, it specifies the maximum number of original recipients for any message. Additional recipients that are generated by aliasing or forwarding do not count. SMTP messages get a 452 response for all recipients over the limit; earlier recipients are delivered as normal. Non-SMTP messages with too many recipients are failed, and no deliveries are done.

**Note:** The RFCs specify that an SMTP server should accept at least 100 RCPT commands in a single message.

<b>recipients_max_reject</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------------	------------------	----------------------	-----------------------

If this option is set true, Exim rejects SMTP messages containing too many recipients by giving 552 errors to the surplus RCPT commands, and a 554 error to the eventual DATA command. Otherwise (the default) it gives a 452 error to the surplus RCPT commands and accepts the message on behalf of the initial set of recipients. The remote server should then re-send the message for the remaining recipients at a later time.

<b>remote_max_parallel</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>2</i>
----------------------------	------------------	----------------------	-------------------

This option controls parallel delivery of one message to a number of remote hosts. If the value is less than 2, parallel delivery is disabled, and Exim does all the remote deliveries for a message one by one. Otherwise, if a single message has to be delivered to more than one remote host, or if several copies have to be sent to the same remote host, up to **remote\_max\_parallel** deliveries are done simultaneously. If more than **remote\_max\_parallel** deliveries are required, the maximum number of processes are started, and as each one finishes, another is begun. The order of starting processes is the same as if sequential delivery were being done, and can be controlled by the **remote\_sort\_domains** option. If parallel delivery takes place while running with debugging turned on, the debugging output from each delivery process is tagged with its process id.

This option controls only the maximum number of parallel deliveries for one message in one Exim delivery process. Because Exim has no central queue manager, there is no way of controlling the total number of simultaneous deliveries if the configuration allows a delivery attempt as soon as a message is received.

If you want to control the total number of deliveries on the system, you need to set the **queue\_only** option. This ensures that all incoming messages are added to the queue without starting a delivery process. Then set up an Exim daemon to start queue runner processes at appropriate intervals (probably fairly often, for example, every minute), and limit the total number of queue runners by setting the **queue\_run\_max** parameter. Because each queue runner delivers only one message at a time, the maximum number of deliveries that can then take place at once is **queue\_run\_max** multiplied by **remote\_max\_parallel**.

If it is purely remote deliveries you want to control, use **queue\_smtp\_domains** instead of **queue\_only**. This has the added benefit of doing the SMTP routing before queueing, so that several messages for the same host will eventually get delivered down the same connection.

<b>remote_sort_domains</b>	Use: <i>main</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
----------------------------	------------------	---------------------------	-----------------------

When there are a number of remote deliveries for a message, they are sorted by domain into the order given by this list. For example,

```
remote_sort_domains = *.cam.ac.uk:*.uk
```

would attempt to deliver to all addresses in the *cam.ac.uk* domain first, then to those in the **uk** domain, then to any others.

<b>retry_data_expire</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>7d</i>
--------------------------	------------------	-------------------	--------------------

This option sets a “use before” time on retry information in Exim’s hints database. Any older retry data is ignored. This means that, for example, once a host has not been tried for 7 days, Exim behaves as if it has no knowledge of past failures.

<b>retry_interval_max</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>24h</i>
---------------------------	------------------	-------------------	---------------------

Chapter 32 describes Exim’s mechanisms for controlling the intervals between delivery attempts for messages that cannot be delivered straight away. This option sets an overall limit to the length of time between retries. It cannot be set greater than 24 hours; any attempt to do so forces the default value.

<b>return_path_remove</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------------------	------------------	----------------------	----------------------

RFC 2821, section 4.4, states that an SMTP server must insert a *Return-path:* header line into a message when it makes a “final delivery”. The *Return-path:* header preserves the sender address as received in the MAIL command. This description implies that this header should not be present in an incoming message. If **return\_path\_remove** is true, any existing *Return-path:* headers are removed from messages at the time they are received. Exim’s transports have options for adding *Return-path:* headers at the time of delivery. They are normally used only for final local deliveries.

<b>return_size_limit</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>100K</i>
--------------------------	------------------	----------------------	----------------------

This option is an obsolete synonym for **bounce\_return\_size\_limit**.

<b>rfc1413_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>@[]</i>
----------------------	------------------	-------------------------	---------------------

RFC 1413 identification calls are made to any client host which matches an item in the list. The default value specifies just this host, being any local interface for the system.

<b>rfc1413_query_timeout</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>0s</i>
------------------------------	------------------	-------------------	--------------------

This sets the timeout on RFC 1413 identification calls. If it is set to zero, no RFC 1413 calls are ever made.

<b>sender_unqualified_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>unset</i>
---------------------------------	------------------	-------------------------	-----------------------

This option lists those hosts from which Exim is prepared to accept unqualified sender addresses. The addresses are made fully qualified by the addition of **qualify\_domain**. This option also affects message header lines. Exim does not reject unqualified addresses in headers that contain sender addresses, but it qualifies them only if the message came from a host that matches **sender\_unqualified\_hosts**, or if the message was submitted locally (not using TCP/IP), and the **-bnq** option was not set.

<b>set_environment</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>empty</i>
------------------------	------------------	--------------------------	-----------------------

This option allows to set individual environment variables that the currently linked libraries and programs in child processes use. The default list is empty,

<b>slow_lookup_log</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>0</i>
------------------------	------------------	----------------------	-------------------

This option controls logging of slow lookups. If the value is nonzero it is taken as a number of milliseconds and lookups taking longer than this are logged. Currently this applies only to DNS lookups.

<b>smtp_accept_keepalive</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
------------------------------	------------------	----------------------	----------------------

This option controls the setting of the SO\_KEEPALIVE option on incoming TCP/IP socket connections. When set, it causes the kernel to probe idle connections periodically, by sending packets with “old” sequence numbers. The other end of the connection should send an acknowledgment if the connection is still okay or a reset if the connection has been aborted. The reason for doing this is that it has the beneficial effect of freeing up certain types of connection that can get stuck when the remote host is disconnected without tidying up the TCP/IP call properly. The keepalive mechanism takes several hours to detect unreachable hosts.

<b>smtp_accept_max</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>20</i>
------------------------	------------------	----------------------	--------------------

This option specifies the maximum number of simultaneous incoming SMTP calls that Exim will accept. It applies only to the listening daemon; there is no control (in Exim) when incoming SMTP is being handled by *inetd*. If the value is set to zero, no limit is applied. However, it is required to be non-zero if either **smtp\_accept\_max\_per\_host** or **smtp\_accept\_queue** is set. See also **smtp\_accept\_reserve** and **smtp\_load\_reserve**.

A new SMTP connection is immediately rejected if the **smtp\_accept\_max** limit has been reached. If not, Exim first checks **smtp\_accept\_max\_per\_host**. If that limit has not been reached for the client host, **smtp\_accept\_reserve** and **smtp\_load\_reserve** are then checked before accepting the connection.

<b>smtp_accept_max_nonmail</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>10</i>
--------------------------------	------------------	----------------------	--------------------

Exim counts the number of “non-mail” commands in an SMTP session, and drops the connection if there are too many. This option defines “too many”. The check catches some denial-of-service attacks, repeated failing AUTHs, or a mad client looping sending EHLO, for example. The check is applied only if the client host matches **smtp\_accept\_max\_nonmail\_hosts**.

When a new message is expected, one occurrence of RSET is not counted. This allows a client to send one RSET between messages (this is not necessary, but some clients do it). Exim also allows one uncounted occurrence of HELO or EHLO, and one occurrence of STARTTLS between messages. After starting up a TLS session, another EHLO is expected, and so it too is not counted. The first occurrence of AUTH in a connection, or immediately following STARTTLS is not counted. Otherwise, all commands other than MAIL, RCPT, DATA, and QUIT are counted.

<b>smtp_accept_max_nonmail_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>*</i>
--------------------------------------	------------------	-------------------------	-------------------

You can control which hosts are subject to the **smtp\_accept\_max\_nonmail** check by setting this option. The default value makes it apply to all hosts. By changing the value, you can exclude any badly-behaved hosts that you have to live with.

<b>smtp_accept_max_per_connection</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>1000</i>
---------------------------------------	------------------	----------------------	----------------------

The value of this option limits the number of MAIL commands that Exim is prepared to accept over a single SMTP connection, whether or not each command results in the transfer of a message. After the limit is reached, a 421 response is given to subsequent MAIL commands. This limit is a safety precaution against a client that goes mad (incidents of this type have been seen).

<b>smtp_accept_max_per_host</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
---------------------------------	------------------	----------------------	-----------------------

This option restricts the number of simultaneous IP connections from a single host (strictly, from a single IP address) to the Exim daemon. The option is expanded, to enable different limits to be applied to different hosts by reference to *\$sender\_host\_address*. Once the limit is reached, additional connection attempts from the same host are rejected with error code 421. This is entirely independent of **smtp\_accept\_reserve**. The option's default value of zero imposes no limit. If this option is set greater than zero, it is required that **smtp\_accept\_max** be non-zero.

**Warning:** When setting this option you should not use any expansion constructions that take an appreciable amount of time. The expansion and test happen in the main daemon loop, in order to reject additional connections without forking additional processes (otherwise a denial-of-service attack could cause a vast number of processes to be created). While the daemon is doing this processing, it cannot accept any other incoming connections.

<b>smtp_accept_queue</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>0</i>
--------------------------	------------------	----------------------	-------------------

If the number of simultaneous incoming SMTP connections being handled via the listening daemon exceeds this value, messages received by SMTP are just placed on the queue; no delivery processes are started automatically. The count is fixed at the start of an SMTP connection. It cannot be updated in the subprocess that receives messages, and so the queueing or not queueing applies to all messages received in the same connection.

A value of zero implies no limit, and clearly any non-zero value is useful only if it is less than the **smtp\_accept\_max** value (unless that is zero). See also **queue\_only**, **queue\_only\_load**, **queue\_smtp\_domains**, and the various **-odx** command line options.

<b>smtp_accept_queue_per_connection</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>10</i>
---	------------------	----------------------	--------------------

This option limits the number of delivery processes that Exim starts automatically when receiving messages via SMTP, whether via the daemon or by the use of **-bs** or **-bS**. If the value of the option is greater than zero, and the number of messages received in a single SMTP session exceeds this number, subsequent messages are placed on the queue, but no delivery processes are started. This helps to limit the number of Exim processes when a server restarts after downtime and there is a lot of mail waiting for it on other systems. On large systems, the default should probably be increased, and on dial-in client systems it should probably be set to zero (that is, disabled).

<b>smtp_accept_reserve</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>0</i>
----------------------------	------------------	----------------------	-------------------

When **smtp\_accept\_max** is set greater than zero, this option specifies a number of SMTP connections that are reserved for connections from the hosts that are specified in **smtp\_reserve\_hosts**. The value set in **smtp\_accept\_max** includes this reserve pool. The specified hosts are not restricted to this number of connections; the option specifies a minimum number of connection slots for them, not a maximum. It is a guarantee that this group of hosts can always get at least **smtp\_accept\_reserve** connections. However, the limit specified by **smtp\_accept\_max\_per\_host** is still applied to each individual host.

For example, if **smtp\_accept\_max** is set to 50 and **smtp\_accept\_reserve** is set to 5, once there are 45 active connections (from any hosts), new connections are accepted only from hosts listed in **smtp\_reserve\_hosts**, provided the other criteria for acceptance are met.

<b>smtp_active_hostname</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
-----------------------------	------------------	----------------------	-----------------------

This option is provided for multi-homed servers that want to masquerade as several different hosts. At the start of an incoming SMTP connection, its value is expanded and used instead of the value of *\$primary\_hostname* in SMTP responses. For example, it is used as domain name in the response to an incoming HELO or EHLO command.

The active hostname is placed in the *\$smtp\_active\_hostname* variable, which is saved with any messages that are received. It is therefore available for use in routers and transports when the message is later delivered.

If this option is unset, or if its expansion is forced to fail, or if the expansion results in an empty string, the value of *\$primary\_hostname* is used. Other expansion failures cause a message to be written to the main and panic logs, and the SMTP command receives a temporary error. Typically, the value of **smtp\_active\_hostname** depends on the incoming interface address. For example:

```
smtp_active_hostname = ${if eq{$received_ip_address}{10.0.0.1}\
    {cox.mydomain}{box.mydomain}}
```

Although *\$smtp\_active\_hostname* is primarily concerned with incoming messages, it is also used as the default for HELO commands in callout verification if there is no remote transport from which to obtain a **helo\_data** value.

<b>smtp_banner</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>see below</i>
--------------------	------------------	----------------------	---------------------------

This string, which is expanded every time it is used, is output as the initial positive response to an SMTP connection. The default setting is:

```
smtp_banner = $smtp_active_hostname ESMTP Exim \
    $version_number $tod_full
```

Failure to expand the string causes a panic error. If you want to create a multiline response to the initial SMTP connection, use “\n” in the string at appropriate points, but not at the end. Note that the 220 code is not included in this string. Exim adds it automatically (several times in the case of a multiline response).

<b>smtp_check_spool_space</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
-------------------------------	------------------	----------------------	----------------------

When this option is set, if an incoming SMTP session encounters the SIZE option on a MAIL command, it checks that there is enough space in the spool directory’s partition to accept a message of that size, while still leaving free the amount specified by **check\_spool\_space** (even if that value is zero). If there isn’t enough space, a temporary error code is returned.

<b>smtp_connect_backlog</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: <i>20</i>
-----------------------------	------------------	----------------------	--------------------

This option specifies a maximum number of waiting SMTP connections. Exim passes this value to the TCP/IP system when it sets up its listener. Once this number of connections are waiting for the daemon's attention, subsequent connection attempts are refused at the TCP/IP level. At least, that is what the manuals say; in some circumstances such connection attempts have been observed to time out instead. For large systems it is probably a good idea to increase the value (to 50, say). It also gives some protection against denial-of-service attacks by SYN flooding.

<b>smtp_enforce_sync</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
--------------------------	------------------	----------------------	----------------------

The SMTP protocol specification requires the client to wait for a response from the server at certain points in the dialogue. Without PIPELINING these synchronization points are after every command; with PIPELINING they are fewer, but they still exist.

Some spamming sites send out a complete set of SMTP commands without waiting for any response. Exim protects against this by rejecting a message if the client has sent further input when it should not have. The error response "554 SMTP synchronization error" is sent, and the connection is dropped. Testing for this error cannot be perfect because of transmission delays (unexpected input may be on its way but not yet received when Exim checks). However, it does detect many instances.

The check can be globally disabled by setting **smtp\_enforce\_sync** false. If you want to disable the check selectively (for example, only for certain hosts), you can do so by an appropriate use of a **control** modifier in an ACL (see section 43.22). See also **pipelining\_advertise\_hosts**.

<b>smtp_etrn_command</b>	Use: <i>main</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
--------------------------	------------------	----------------------------------	-----------------------

If this option is set, the given command is run whenever an SMTP ETRN command is received from a host that is permitted to issue such commands (see chapter 43). The string is split up into separate arguments which are independently expanded. The expansion variable *\$domain* is set to the argument of the ETRN command, and no syntax checking is done on it. For example:

```
smtp_etrn_command = /etc/etrn_command $domain \
                    $sender_host_address
```

A new process is created to run the command, but Exim does not wait for it to complete. Consequently, its status cannot be checked. If the command cannot be run, a line is written to the panic log, but the ETRN caller still receives a 250 success response. Exim is normally running under its own uid when receiving SMTP, so it is not possible for it to change the uid before running the command.

<b>smtp_etrn_serialize</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
----------------------------	------------------	----------------------	----------------------

When this option is set, it prevents the simultaneous execution of more than one identical command as a result of ETRN in an SMTP connection. See section 48.8 for details.

<b>smtp_load_reserve</b>	Use: <i>main</i>	Type: <i>fixed-point</i>	Default: <i>unset</i>
--------------------------	------------------	--------------------------	-----------------------

If the system load average ever gets higher than this, incoming SMTP calls are accepted only from those hosts that match an entry in **smtp\_reserve\_hosts**. If **smtp\_reserve\_hosts** is not set, no incoming SMTP calls are accepted when the load is over the limit. The option has no effect on ancient operating systems on which Exim cannot determine the load average. See also **deliver\_queue\_load\_max** and **queue\_only\_load**.



<b>smtp_max_synprot_errors</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: 3
--------------------------------	------------------	----------------------	------------

Exim rejects SMTP commands that contain syntax or protocol errors. In particular, a syntactically invalid email address, as in this command:

```
RCPT TO:<abc xyz@a.b.c>
```

causes immediate rejection of the command, before any other tests are done. (The ACL cannot be run if there is no valid address to set up for it.) An example of a protocol error is receiving RCPT before MAIL. If there are too many syntax or protocol errors in one SMTP session, the connection is dropped. The limit is set by this option.

When the PIPELINING extension to SMTP is in use, some protocol errors are “expected”, for instance, a RCPT command after a rejected MAIL command. Exim assumes that PIPELINING will be used if it advertises it (see **pipelining\_advertise\_hosts**), and in this situation, “expected” errors do not count towards the limit.

<b>smtp_max_unknown_commands</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: 3
----------------------------------	------------------	----------------------	------------

If there are too many unrecognized commands in an incoming SMTP session, an Exim server drops the connection. This is a defence against some kinds of abuse that subvert web clients into making connections to SMTP ports; in these circumstances, a number of non-SMTP command lines are sent first.

<b>smtp_ratelimit_hosts</b>	Use: <i>main</i>	Type: <i>host list</i> <sup>†</sup>	Default: <i>unset</i>
-----------------------------	------------------	-------------------------------------	-----------------------

Some sites find it helpful to be able to limit the rate at which certain hosts can send them messages, and the rate at which an individual message can specify recipients.

Exim has two rate-limiting facilities. This section describes the older facility, which can limit rates within a single connection. The newer **ratelimit** ACL condition can limit rates across all connections. See section 43.38 for details of the newer facility.

When a host matches **smtp\_ratelimit\_hosts**, the values of **smtp\_ratelimit\_mail** and **smtp\_ratelimit\_rcpt** are used to control the rate of acceptance of MAIL and RCPT commands in a single SMTP session, respectively. Each option, if set, must contain a set of four comma-separated values:

- A threshold, before which there is no rate limiting.
- An initial time delay. Unlike other times in Exim, numbers with decimal fractional parts are allowed here.
- A factor by which to increase the delay each time.
- A maximum value for the delay. This should normally be less than 5 minutes, because after that time, the client is liable to timeout the SMTP command.

For example, these settings have been used successfully at the site which first suggested this feature, for controlling mail from their customers:

```
smtp_ratelimit_mail = 2,0.5s,1.05,4m
smtp_ratelimit_rcpt = 4,0.25s,1.015,4m
```

The first setting specifies delays that are applied to MAIL commands after two have been received over a single connection. The initial delay is 0.5 seconds, increasing by a factor of 1.05 each time. The second setting applies delays to RCPT commands when more than four occur in a single message.

<b>smtp_ratelimit_mail</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------------	------------------	---------------------	-----------------------

See **smtp\_ratelimit\_hosts** above.

<b>smtp_ratelimit_rcpt</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------------	------------------	---------------------	-----------------------

See **smtp\_ratelimit\_hosts** above.

<b>smtp_receive_timeout</b>	Use: <i>main</i>	Type: <i>time</i> †	Default: <i>5m</i>
-----------------------------	------------------	---------------------	--------------------

This sets a timeout value for SMTP reception. It applies to all forms of SMTP input, including batch SMTP. If a line of input (either an SMTP command or a data line) is not received within this time, the SMTP connection is dropped and the message is abandoned. A line is written to the log containing one of the following messages:

```
SMTP command timeout on connection from...
SMTP data timeout on connection from...
```

The former means that Exim was expecting to read an SMTP command; the latter means that it was in the DATA phase, reading the contents of a message.

If the first character of the option is a “\$” the option is expanded before use and may depend on *\$sender\_host\_name*, *\$sender\_host\_address* and *\$sender\_host\_port*.

The value set by this option can be overridden by the **-os** command-line option. A setting of zero time disables the timeout, but this should never be used for SMTP over TCP/IP. (It can be useful in some cases of local input using **-bs** or **-bS**.) For non-SMTP input, the reception timeout is controlled by **receive\_timeout** and **-or**.

<b>smtp_reserve_hosts</b>	Use: <i>main</i>	Type: <i>host list</i> †	Default: <i>unset</i>
---------------------------	------------------	--------------------------	-----------------------

This option defines hosts for which SMTP connections are reserved; see **smtp\_accept\_reserve** and **smtp\_load\_reserve** above.

<b>smtp_return_error_details</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------------------	------------------	----------------------	-----------------------

In the default state, Exim uses bland messages such as “Administrative prohibition” when it rejects SMTP commands for policy reasons. Many sysadmins like this because it gives away little information to spammers. However, some other sysadmins who are applying strict checking policies want to give out much fuller information about failures. Setting **smtp\_return\_error\_details** true causes Exim to be more forthcoming. For example, instead of “Administrative prohibition”, it might give:

```
550-Rejected after DATA: '>' missing at end of address:
550 failing address in "From" header is: <user@dom.ain
```

<b>smtputf8_advertise_hosts</b>	Use: <i>main</i>	Type: <i>host list</i> †	Default: <i>*</i>
---------------------------------	------------------	--------------------------	-------------------

When Exim is built with support for internationalised mail names, the availability thereof is advertised in response to EHLO only to those client hosts that match this option. See chapter 59 for details of Exim’s support for internationalisation.

<b>spamd_address</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>see below</i>
----------------------	------------------	---------------------	---------------------------

This option is available when Exim is compiled with the content-scanning extension. It specifies how Exim connects to SpamAssassin’s **spamd** daemon. The default value is

```
127.0.0.1 783
```

See section 44.2 for more details.

<b>split_spool_directory</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------------	------------------	----------------------	-----------------------

If this option is set, it causes Exim to split its input directory into 62 subdirectories, each with a single alphanumeric character as its name. The sixth character of the message id is used to allocate messages to subdirectories; this is the least significant base-62 digit of the time of arrival of the message.

Splitting up the spool in this way may provide better performance on systems where there are long mail queues, by reducing the number of files in any one directory. The msglog directory is also split up in a similar way to the input directory; however, if **preserve\_message\_logs** is set, all old msglog files are still placed in the single directory *msglog.OLD*.

It is not necessary to take any special action for existing messages when changing **split\_spool\_directory**. Exim notices messages that are in the “wrong” place, and continues to process them. If the option is turned off after a period of being on, the subdirectories will eventually empty and be automatically deleted.

When **split\_spool\_directory** is set, the behaviour of queue runner processes changes. Instead of creating a list of all messages in the queue, and then trying to deliver each one in turn, it constructs a list of those in one sub-directory and tries to deliver them, before moving on to the next sub-directory. The sub-directories are processed in a random order. This spreads out the scanning of the input directories, and uses less memory. It is particularly beneficial when there are lots of messages on the queue. However, if **queue\_run\_in\_order** is set, none of this new processing happens. The entire queue has to be scanned and sorted before any deliveries can start.

<b>spool_directory</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>set at compile time</i>
------------------------	------------------	----------------------	-------------------------------------

This defines the directory in which Exim keeps its spool, that is, the messages it is waiting to deliver. The default value is taken from the compile-time configuration setting, if there is one. If not, this option must be set. The string is expanded, so it can contain, for example, a reference to *\$primary\_hostname*.

If the spool directory name is fixed on your installation, it is recommended that you set it at build time rather than from this option, particularly if the log files are being written to the spool directory (see **log\_file\_path**). Otherwise log files cannot be used for errors that are detected early on, such as failures in the configuration file.

By using this option to override the compiled-in path, it is possible to run tests of Exim without using the standard spool.

<b>spool_wireformat</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------	------------------	----------------------	-----------------------

If this option is set, Exim may for some messages use an alternate format for data-files in the spool which matches the wire format. Doing this permits more efficient message reception and transmission. Currently it is only done for messages received using the EMSTP CHUNKING option.

The following variables will not have useful values:

```
$max_received_linelength
$body_linecount
$body_zerocount
```

Users of the local\_scan() API (see 45), and any external programs which are passed a reference to a message data file (except via the “regex”, “malware” or “spam”) ACL conditions) will need to be aware of the potential different format.

Using any of the ACL conditions noted will negate the reception benefit (as a Unix-mbox-format file is constructed for them). The transmission benefit is maintained.

<b>sqlite_lock_timeout</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>5s</i>
----------------------------	------------------	-------------------	--------------------

This option controls the timeout that the *sqlite* lookup uses when trying to access an SQLite database. See section 9.26 for more details.

<b>strict_acl_vars</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------	----------------------	-----------------------

This option controls what happens if a syntactically valid but undefined ACL variable is referenced. If it is false (the default), an empty string is substituted; if it is true, an error is generated. See section 43.19 for details of ACL variables.

<b>strip_excess_angle_brackets</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------------------	------------------	----------------------	-----------------------

If this option is set, redundant pairs of angle brackets round “route-addr” items in addresses are stripped. For example, <<xxx@a.b.c.d>> is treated as <xxx@a.b.c.d>. If this is in the envelope and the message is passed on to another MTA, the excess angle brackets are not passed on. If this option is not set, multiple pairs of angle brackets cause a syntax error.

<b>strip_trailing_dot</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	------------------	----------------------	-----------------------

If this option is set, a trailing dot at the end of a domain in an address is ignored. If this is in the envelope and the message is passed on to another MTA, the dot is not passed on. If this option is not set, a dot at the end of a domain causes a syntax error. However, addresses in header lines are checked only when an ACL requests header syntax checking.

<b>syslog_duplication</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------------------	------------------	----------------------	----------------------

When Exim is logging to syslog, it writes the log lines for its three separate logs at different syslog priorities so that they can in principle be separated on the logging hosts. Some installations do not require this separation, and in those cases, the duplication of certain log lines is a nuisance. If **syslog\_duplication** is set false, only one copy of any particular log line is written to syslog. For lines that normally go to both the main log and the reject log, the reject log version (possibly containing message header lines) is written, at LOG\_NOTICE priority. Lines that normally go to both the main and the panic log are written at the LOG\_ALERT priority.

<b>syslog_facility</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
------------------------	------------------	---------------------	-----------------------

This option sets the syslog “facility” name, used when Exim is logging to syslog. The value must be one of the strings “mail”, “user”, “news”, “uucp”, “daemon”, or “localx” where *x* is a digit between 0 and 7. If this option is unset, “mail” is used. See chapter 52 for details of Exim’s logging.

<b>syslog_pid</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
-------------------	------------------	----------------------	----------------------

If **syslog\_pid** is set false, the PID on Exim’s log lines are omitted when these lines are sent to syslog. (Syslog normally prefixes the log lines with the PID of the logging process automatically.) You need to enable the +pid log selector item, if you want Exim to write it’s PID into the logs.) See chapter 52 for details of Exim’s logging.

<b>syslog_processname</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>exim</i>
---------------------------	------------------	---------------------	----------------------

This option sets the syslog “ident” name, used when Exim is logging to syslog. The value must be no longer than 32 characters. See chapter 52 for details of Exim’s logging.

<b>syslog_timestamp</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
-------------------------	------------------	----------------------	----------------------

If **syslog\_timestamp** is set false, the timestamps on Exim's log lines are omitted when these lines are sent to syslog. See chapter 52 for details of Exim's logging.

<b>system_filter</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option specifies an Exim filter file that is applied to all messages at the start of each delivery attempt, before any routing is done. System filters must be Exim filters; they cannot be Sieve filters. If the system filter generates any deliveries to files or pipes, or any new mail messages, the appropriate **system\_filter\_...\_transport** option(s) must be set, to define which transports are to be used. Details of this facility are given in chapter 46. A forced expansion failure results in no filter operation.

<b>system_filter_directory_transport</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
--	------------------	----------------------	-----------------------

This sets the name of the transport driver that is to be used when the **save** command in a system message filter specifies a path ending in “/”, implying delivery of each message into a separate file in some directory. During the delivery, the variable *\$address\_file* contains the path name.

<b>system_filter_file_transport</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------------------	------------------	----------------------	-----------------------

This sets the name of the transport driver that is to be used when the **save** command in a system message filter specifies a path not ending in “/”. During the delivery, the variable *\$address\_file* contains the path name.

<b>system_filter_group</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------------	------------------	---------------------	-----------------------

This option is used only when **system\_filter\_user** is also set. It sets the gid under which the system filter is run, overriding any gid that is associated with the user. The value may be numerical or symbolic.

<b>system_filter_pipe_transport</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------------------	------------------	----------------------	-----------------------

This specifies the transport driver that is to be used when a **pipe** command is used in a system filter. During the delivery, the variable *\$address\_pipe* contains the pipe command.

<b>system_filter_reply_transport</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------------------------	------------------	----------------------	-----------------------

This specifies the transport driver that is to be used when a **mail** command is used in a system filter.

<b>system_filter_user</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
---------------------------	------------------	---------------------	-----------------------

If this option is set to root, the system filter is run in the main Exim delivery process, as root. Otherwise, the system filter runs in a separate process, as the given user, defaulting to the Exim run-time user. Unless the string consists entirely of digits, it is looked up in the password data. Failure to find the named user causes a configuration error. The gid is either taken from the password data, or specified by **system\_filter\_group**. When the uid is specified numerically, **system\_filter\_group** is required to be set.

If the system filter generates any pipe, file, or reply deliveries, the uid under which the filter is run is used when transporting them, unless a transport option overrides.

<b>tcp_nodelay</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
--------------------	------------------	----------------------	----------------------

If this option is set false, it stops the Exim daemon setting the TCP\_NODELAY option on its listening sockets. Setting TCP\_NODELAY turns off the “Nagle algorithm”, which is a way of improving network performance in interactive (character-by-character) situations. Turning it off should improve Exim’s performance a bit, so that is what happens by default. However, it appears that some broken clients cannot cope, and time out. Hence this option. It affects only those sockets that are set up for listening by the daemon. Sockets created by the smtp transport for delivering mail always set TCP\_NODELAY.

<b>timeout_frozen_after</b>	Use: <i>main</i>	Type: <i>time</i>	Default: <i>0s</i>
-----------------------------	------------------	-------------------	--------------------

If **timeout\_frozen\_after** is set to a time greater than zero, a frozen message of any kind that has been on the queue for longer than the given time is automatically cancelled at the next queue run. If the frozen message is a bounce message, it is just discarded; otherwise, a bounce is sent to the sender, in a similar manner to cancellation by the **-Mg** command line option. If you want to timeout frozen bounce messages earlier than other kinds of frozen message, see **ignore\_bounce\_errors\_after**.

**Note:** the default value of zero means no timeouts; with this setting, frozen messages remain on the queue forever (except for any frozen bounce messages that are released by **ignore\_bounce\_errors\_after**).

<b>timezone</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
-----------------	------------------	---------------------	-----------------------

The value of **timezone** is used to set the environment variable TZ while running Exim (if it is different on entry). This ensures that all timestamps created by Exim are in the required timezone. If you want all your timestamps to be in UTC (aka GMT) you should set

```
timezone = UTC
```

The default value is taken from TIMEZONE\_DEFAULT in *Local/Makefile*, or, if that is not set, from the value of the TZ environment variable when Exim is built. If **timezone** is set to the empty string, either at build or run time, any existing TZ variable is removed from the environment when Exim runs. This is appropriate behaviour for obtaining wall-clock time on some, but unfortunately not all, operating systems.

<b>tls_advertise_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>*</i>
----------------------------	------------------	-------------------------	-------------------

When Exim is built with support for TLS encrypted connections, the availability of the STARTTLS command to set up an encrypted session is advertised in response to EHLO only to those client hosts that match this option. See chapter 42 for details of Exim’s support for TLS. Note that the default value requires that a certificate be supplied using the **tls\_certificate** option. If TLS support for incoming connections is not required the **tls\_advertise\_hosts** option should be set empty.

<b>tls_certificate</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>list†</i>
------------------------	------------------	---------------------	-----------------------

The value of this option is expanded, and must then be a list of absolute paths to files which contains the server’s certificates. Commonly only one file is needed.

The server’s private key is also assumed to be in this file if **tls\_privatekey** is unset. See chapter 42 for further details.

**Note:** The certificates defined by this option are used only when Exim is receiving incoming messages as a server. If you want to supply certificates for use when sending messages as a client, you must set the **tls\_certificate** option in the relevant *smtp* transport.

**Note:** If you use filenames based on IP addresses, change the list separator in the usual way to avoid confusion under IPv6.

**Note:** Under current versions of OpenSSL, when a list of more than one file is used, the `$tls_in_ourcert` variable is unreliable.

**Note:** OCSP stapling is not usable when a list of more than one file is used.

If the option contains `$tls_out_sni` and Exim is built against OpenSSL, then if the OpenSSL build supports TLS extensions and the TLS client sends the Server Name Indication extension, then this option and others documented in 42.10 will be re-expanded.

If this option is unset or empty a fresh self-signed certificate will be generated for every connection.

<b>tls_crl</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------	------------------	----------------------	-----------------------

This option specifies a certificate revocation list. The expanded value must be the name of a file that contains a CRL in PEM format.

See 42.10 for discussion of when this option might be re-expanded.

<b>tls_dh_max_bits</b>	Use: <i>main</i>	Type: <i>integer</i>	Default: 2236
------------------------	------------------	----------------------	---------------

The number of bits used for Diffie-Hellman key-exchange may be suggested by the chosen TLS library. That value might prove to be too high for interoperability. This option provides a maximum clamp on the value suggested, trading off security for interoperability.

The value must be at least 1024.

The value 2236 was chosen because, at time of adding the option, it was the hard-coded maximum value supported by the NSS cryptographic library, as used by Thunderbird, while GnuTLS was suggesting 2432 bits as normal.

If you prefer more security and are willing to break some clients, raise this number.

Note that the value passed to GnuTLS for *\*generating\** a new prime may be a little less than this figure, because GnuTLS is inexact and may produce a larger prime than requested.

<b>tls_dhparam</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------	------------------	----------------------	-----------------------

The value of this option is expanded and indicates the source of DH parameters to be used by Exim.

**Note: The Exim Maintainers strongly recommend using a filename with site-generated local DH parameters**, which has been supported across all versions of Exim. The other specific constants available are a fallback so that even when "unconfigured", Exim can offer Perfect Forward Secrecy in older ciphersuites in TLS.

If **tls\_dhparam** is a filename starting with a `/`, then it names a file from which DH parameters should be loaded. If the file exists, it should hold a PEM-encoded PKCS#3 representation of the DH prime. If the file does not exist, for OpenSSL it is an error. For GnuTLS, Exim will attempt to create the file and fill it with a generated DH prime. For OpenSSL, if the DH bit-count from loading the file is greater than **tls\_dh\_max\_bits** then it will be ignored, and treated as though the **tls\_dhparam** were set to "none".

If this option expands to the string "none", then no DH parameters will be loaded by Exim.

If this option expands to the string "historic" and Exim is using GnuTLS, then Exim will attempt to load a file from inside the spool directory. If the file does not exist, Exim will attempt to create it. See section 42.3 for further details.

If Exim is using OpenSSL and this option is empty or unset, then Exim will load a default DH prime; the default is Exim-specific but lacks verifiable provenance.

In older versions of Exim the default was the 2048 bit prime described in section 2.2 of RFC 5114, "2048-bit MODP Group with 224-bit Prime Order Subgroup", which in IKE is assigned number 23.

Otherwise, the option must expand to the name used by Exim for any of a number of DH primes specified in RFC 2409, RFC 3526, RFC 5114, RFC 7919, or from other sources. As names, Exim uses a standard specified name, else "ike" followed by the number used by IKE, or "default" which corresponds to `exim.dev.20160529.3`.

The available standard primes are: `ffdhe2048`, `ffdhe3072`, `ffdhe4096`, `ffdhe6144`, `ffdhe8192`, `ike1`, `ike2`, `ike5`, `ike14`, `ike15`, `ike16`, `ike17`, `ike18`, `ike22`, `ike23` and `ike24`.

The available additional primes are: `exim.dev.20160529.1`, `exim.dev.20160529.2` and `exim.dev.20160529.3`.

Some of these will be too small to be accepted by clients. Some may be too large to be accepted by clients. The open cryptographic community has suspicions about the integrity of some of the later IKE values, which led into RFC7919 providing new fixed constants (the "ffdhe" identifiers).

At this point, all of the "ike" values should be considered obsolete; they're still in Exim to avoid breaking unusual configurations, but are candidates for removal the next time we have backwards-incompatible changes.

The TLS protocol does not negotiate an acceptable size for this; clients tend to hard-drop connections if what is offered by the server is unacceptable, whether too large or too small, and there's no provision for the client to tell the server what these constraints are. Thus, as a server operator, you need to make an educated guess as to what is most likely to work for your userbase.

Some known size constraints suggest that a bit-size in the range 2048 to 2236 is most likely to maximise interoperability. The upper bound comes from applications using the Mozilla Network Security Services (NSS) library, which used to set its `DH_MAX_P_BITS` upper-bound to 2236. This affects many mail user agents (MUAs). The lower bound comes from Debian installs of Exim4 prior to the 4.80 release, as Debian used to patch Exim to raise the minimum acceptable bound from 1024 to 2048.

<b>tls_eccurve</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>auto</i>
--------------------	------------------	----------------------	----------------------

This option selects a EC curve for use by Exim when used with OpenSSL. It has no effect when Exim is used with GnuTLS.

After expansion it must contain a valid EC curve parameter, such as `prime256v1`, `secp384r1`, or `P-512`. Consult your OpenSSL manual for valid selections.

For OpenSSL versions before (and not including) 1.0.2, the string `auto` selects `prime256v1`. For more recent OpenSSL versions `auto` tells the library to choose.

If the option expands to an empty string, no EC curves will be enabled.

<b>tls_ocsp_file</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This option must if set expand to the absolute path to a file which contains a current status proof for the server's certificate, as obtained from the Certificate Authority.

Usable for GnuTLS 3.4.4 or 3.3.17 or OpenSSL 1.1.0 (or later).

**Note:** There is currently no support for multiple OCSP proofs to match the multiple certificates facility.

<b>tls_on_connect_ports</b>	Use: <i>main</i>	Type: <i>string list</i>	Default: <i>unset</i>
-----------------------------	------------------	--------------------------	-----------------------

This option specifies a list of incoming SSMTP (aka SMTPS) ports that should operate the obsolete SSMTP (SMTPS) protocol, where a TLS session is immediately set up without waiting for the client to issue a STARTTLS command. For further details, see section 13.4.



<b>tls_privatekey</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>list†</i>
-----------------------	------------------	---------------------	-----------------------

The value of this option is expanded, and must then be a list of absolute paths to files which contains the server's private keys.

If this option is unset, or if the expansion is forced to fail, or the result is an empty string, the private key is assumed to be in the same file as the server's certificates. See chapter 42 for further details.

See 42.10 for discussion of when this option might be re-expanded.

<b>tls_remember_esmtp</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	------------------	----------------------	-----------------------

If this option is set true, Exim violates the RFCs by remembering that it is in "esmtp" state after successfully negotiating a TLS session. This provides support for broken clients that fail to send a new EHLO after starting a TLS session.

<b>tls_require_ciphers</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------------	------------------	----------------------	-----------------------

This option controls which ciphers can be used for incoming TLS connections. The *smtp* transport has an option of the same name for controlling outgoing connections. This option is expanded for each connection, so can be varied for different clients if required. The value of this option must be a list of permitted cipher suites. The OpenSSL and GnuTLS libraries handle cipher control in somewhat different ways. If GnuTLS is being used, the client controls the preference order of the available ciphers. Details are given in sections 42.4 and 42.5.

<b>tls_try_verify_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>unset</i>
-----------------------------	------------------	-------------------------	-----------------------

See **tls\_verify\_hosts** below.

<b>tls_verify_certificates</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>system</i>
--------------------------------	------------------	----------------------	------------------------

The value of this option is expanded, and must then be either the word "system" or the absolute path to a file or directory containing permitted certificates for clients that match **tls\_verify\_hosts** or **tls\_try\_verify\_hosts**.

The "system" value for the option will use a system default location compiled into the SSL library. This is not available for GnuTLS versions preceding 3.0.20, and will be taken as empty; an explicit location must be specified.

The use of a directory for the option value is not available for GnuTLS versions preceding 3.3.6 and a single file must be used.

With OpenSSL the certificates specified explicitly either by file or directory are added to those given by the system default location.

These certificates should be for the certificate authorities trusted, rather than the public cert of individual clients. With both OpenSSL and GnuTLS, if the value is a file then the certificates are sent by Exim as a server to connecting clients, defining the list of accepted certificate authorities. Thus the values defined should be considered public data. To avoid this, use the explicit directory version.

See 42.10 for discussion of when this option might be re-expanded.

A forced expansion failure or setting to an empty string is equivalent to being unset.

<b>tls_verify_hosts</b>	Use: <i>main</i>	Type: <i>host list†</i>	Default: <i>unset</i>
-------------------------	------------------	-------------------------	-----------------------

This option, along with **tls\_try\_verify\_hosts**, controls the checking of certificates from clients. The expected certificates are defined by **tls\_verify\_certificates**, which must be set. A configuration error occurs if either **tls\_verify\_hosts** or **tls\_try\_verify\_hosts** is set and **tls\_verify\_certificates** is not set.

Any client that matches **tls\_verify\_hosts** is constrained by **tls\_verify\_certificates**. When the client initiates a TLS session, it must present one of the listed certificates. If it does not, the connection is aborted. **Warning:** Including a host in **tls\_verify\_hosts** does not require the host to use TLS. It can still send SMTP commands through unencrypted connections. Forcing a client to use TLS has to be done separately using an ACL to reject inappropriate commands when the connection is not encrypted.

A weaker form of checking is provided by **tls\_try\_verify\_hosts**. If a client matches this option (but not **tls\_verify\_hosts**), Exim requests a certificate and checks it against **tls\_verify\_certificates**, but does not abort the connection if there is no certificate or if it does not match. This state can be detected in an ACL, which makes it possible to implement policies such as “accept for relay only if a verified certificate has been received, but accept for local delivery if encrypted, even without a verified certificate”.

Client hosts that match neither of these lists are not asked to present certificates.

<b>trusted_groups</b>	Use: <i>main</i>	Type: <i>string list†</i>	Default: <i>unset</i>
-----------------------	------------------	---------------------------	-----------------------

This option is expanded just once, at the start of Exim’s processing. If this option is set, any process that is running in one of the listed groups, or which has one of them as a supplementary group, is trusted. The groups can be specified numerically or by name. See section 5.2 for details of what trusted callers are permitted to do. If neither **trusted\_groups** nor **trusted\_users** is set, only root and the Exim user are trusted.

<b>trusted_users</b>	Use: <i>main</i>	Type: <i>string list†</i>	Default: <i>unset</i>
----------------------	------------------	---------------------------	-----------------------

This option is expanded just once, at the start of Exim’s processing. If this option is set, any process that is running as one of the listed users is trusted. The users can be specified numerically or by name. See section 5.2 for details of what trusted callers are permitted to do. If neither **trusted\_groups** nor **trusted\_users** is set, only root and the Exim user are trusted.

<b>unknown_login</b>	Use: <i>main</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

This is a specialized feature for use in unusual configurations. By default, if the uid of the caller of Exim cannot be looked up using *getpwuid()*, Exim gives up. The **unknown\_login** option can be used to set a login name to be used in this circumstance. It is expanded, so values like **user\$caller\_uid** can be set. When **unknown\_login** is used, the value of **unknown\_username** is used for the user’s real name (gecos field), unless this has been set by the **-F** option.

<b>unknown_username</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------------	------------------	---------------------	-----------------------

See **unknown\_login**.

<b>untrusted_set_sender</b>	Use: <i>main</i>	Type: <i>address list†</i>	Default: <i>unset</i>
-----------------------------	------------------	----------------------------	-----------------------

When an untrusted user submits a message to Exim using the standard input, Exim normally creates an envelope sender address from the user’s login and the default qualification domain. Data from the **-f** option (for setting envelope senders on non-SMTP messages) or the SMTP MAIL command (if **-bs** or **-bS** is used) is ignored.

However, untrusted users are permitted to set an empty envelope sender address, to declare that a message should never generate any bounces. For example:

```
exim -f '<>' user@domain.example
```

The **untrusted\_set\_sender** option allows you to permit untrusted users to set other envelope sender addresses in a controlled way. When it is set, untrusted users are allowed to set envelope sender addresses that match any of the patterns in the list. Like all address lists, the string is expanded. The identity of the user is in *\$sender\_ident*, so you can, for example, restrict users to setting senders that start with their login ids followed by a hyphen by a setting like this:

```
untrusted_set_sender = ^$sender_ident-
```

If you want to allow untrusted users to set envelope sender addresses without restriction, you can use

```
untrusted_set_sender = *
```

The **untrusted\_set\_sender** option applies to all forms of local input, but only to the setting of the envelope sender. It does not permit untrusted users to use the other options which trusted user can use to override message parameters. Furthermore, it does not stop Exim from removing an existing *Sender:* header in the message, or from adding a *Sender:* header if necessary. See **local\_sender\_retain** and **local\_from\_check** for ways of overriding these actions. The handling of the *Sender:* header is also described in section 47.16.

The log line for a message's arrival shows the envelope sender following "<=". For local messages, the user's login always follows, after "U=". In **-bp** displays, and in the Exim monitor, if an untrusted user sets an envelope sender address, the user's login is shown in parentheses after the sender address.

<b>uucp_from_pattern</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>see below</i>
--------------------------	------------------	---------------------	---------------------------

Some applications that pass messages to an MTA via a command line interface use an initial line starting with "From " to pass the envelope sender. In particular, this is used by UUCP software. Exim recognizes such a line by means of a regular expression that is set in **uucp\_from\_pattern**. When the pattern matches, the sender address is constructed by expanding the contents of **uucp\_from\_sender**, provided that the caller of Exim is a trusted user. The default pattern recognizes lines in the following two forms:

```
From ph10 Fri Jan 5 12:35 GMT 1996
From ph10 Fri, 7 Jan 97 14:00:00 GMT
```

The pattern can be seen by running

```
exim -bP uucp_from_pattern
```

It checks only up to the hours and minutes, and allows for a 2-digit or 4-digit year in the second case. The first word after "From " is matched in the regular expression by a parenthesized subpattern. The default value for **uucp\_from\_sender** is "\$1", which therefore just uses this first word ("ph10" in the example above) as the message's sender. See also **ignore\_fromline\_hosts**.

<b>uucp_from_sender</b>	Use: <i>main</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>\$1</i>
-------------------------	------------------	----------------------------------	---------------------

See **uucp\_from\_pattern** above.

<b>warn_message_file</b>	Use: <i>main</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------------------	------------------	---------------------	-----------------------

This option defines a template file containing paragraphs of text to be used for constructing the warning message which is sent by Exim when a message has been on the queue for a specified amount of time, as specified by **delay\_warning**. Details of the file's contents are given in chapter 49. See also **bounce\_message\_file**.

<b>write_rejectlog</b>	Use: <i>main</i>	Type: <i>boolean</i>	Default: <i>true</i>
------------------------	------------------	----------------------	----------------------

If this option is set false, Exim no longer writes anything to the reject log. See chapter 52 for details of what Exim writes to its logs.

## 15. Generic options for routers

This chapter describes the generic options that apply to all routers. Those that are preconditions are marked with ‡ in the “use” field.

For a general description of how a router operates, see sections 3.10 and 3.12. The latter specifies the order in which the preconditions are tested. The order of expansion of the options that provide data for a transport is: **errors\_to**, **headers\_add**, **headers\_remove**, **transport**.

<b>address_data</b>	Use: <i>routers</i>	Type: <i>string</i> †	Default: <i>unset</i>
---------------------	---------------------	-----------------------	-----------------------

The string is expanded just before the router is run, that is, after all the precondition tests have succeeded. If the expansion is forced to fail, the router declines, the value of **address\_data** remains unchanged, and the **more** option controls what happens next. Other expansion failures cause delivery of the address to be deferred.

When the expansion succeeds, the value is retained with the address, and can be accessed using the variable *\$address\_data* in the current router, subsequent routers, and the eventual transport.

**Warning:** If the current or any subsequent router is a *redirect* router that runs a user’s filter file, the contents of *\$address\_data* are accessible in the filter. This is not normally a problem, because such data is usually either not confidential or it “belongs” to the current user, but if you do put confidential data into *\$address\_data* you need to remember this point.

Even if the router declines or passes, the value of *\$address\_data* remains with the address, though it can be changed by another **address\_data** setting on a subsequent router. If a router generates child addresses, the value of *\$address\_data* propagates to them. This also applies to the special kind of “child” that is generated by a router with the **unseen** option.

The idea of **address\_data** is that you can use it to look up a lot of data for the address once, and then pick out parts of the data later. For example, you could use a single LDAP lookup to return a string of the form

```
uid=1234 gid=5678 mailbox=/mail/xyz forward=/home/xyz/.forward
```

In the transport you could pick out the mailbox by a setting such as

```
file = ${extract{mailbox}{$address_data}}
```

This makes the configuration file less messy, and also reduces the number of lookups (though Exim does cache lookups).

The **address\_data** facility is also useful as a means of passing information from one router to another, and from a router to a transport. In addition, if *\$address\_data* is set by a router when verifying a recipient address from an ACL, it remains available for use in the rest of the ACL statement. After verifying a sender, the value is transferred to *\$sender\_address\_data*.

<b>address_test</b>	Use: <i>routers</i> ‡	Type: <i>boolean</i>	Default: <i>true</i>
---------------------	-----------------------	----------------------	----------------------

If this option is set false, the router is skipped when routing is being tested by means of the **-bt** command line option. This can be a convenience when your first router sends messages to an external scanner, because it saves you having to set the “already scanned” indicator when testing real address routing.

<b>cannot_route_message</b>	Use: <i>routers</i>	Type: <i>string</i> †	Default: <i>unset</i>
-----------------------------	---------------------	-----------------------	-----------------------

This option specifies a text message that is used when an address cannot be routed because Exim has run out of routers. The default message is “Unrouteable address”. This option is useful only on routers that have **more** set false, or on the very last router in a configuration, because the value that is used is taken from the last router that is considered. This includes a router that is skipped because its

preconditions are not met, as well as a router that declines. For example, using the default configuration, you could put:

```
cannot_route_message = Remote domain not found in DNS
```

on the first router, which is a *dnslookup* router with **more** set false, and

```
cannot_route_message = Unknown local user
```

on the final router that checks for local users. If string expansion fails for this option, the default message is used. Unless the expansion failure was explicitly forced, a message about the failure is written to the main and panic logs, in addition to the normal message about the routing failure.

<b>caseful_local_part</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	---------------------	----------------------	-----------------------

By default, routers handle the local parts of addresses in a case-insensitive manner, though the actual case is preserved for transmission with the message. If you want the case of letters to be significant in a router, you must set this option true. For individual router options that contain address or local part lists (for example, **local\_parts**), case-sensitive matching can be turned on by “+caseful” as a list item. See section 10.20 for more details.

The value of the *\$local\_part* variable is forced to lower case while a router is running unless **caseful\_local\_part** is set. When a router assigns an address to a transport, the value of *\$local\_part* when the transport runs is the same as it was in the router. Similarly, when a router generates child addresses by aliasing or forwarding, the values of *\$original\_local\_part* and *\$parent\_local\_part* are those that were used by the redirecting router.

This option applies to the processing of an address by a router. When a recipient address is being processed in an ACL, there is a separate **control** modifier that can be used to specify case-sensitive processing within the ACL (see section 43.22).

<b>check_local_user</b>	Use: <i>routers</i> ‡	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------	-----------------------	----------------------	-----------------------

When this option is true, Exim checks that the local part of the recipient address (with affixes removed if relevant) is the name of an account on the local system. The check is done by calling the *getpwnam()* function rather than trying to read */etc/passwd* directly. This means that other methods of holding password data (such as NIS) are supported. If the local part is a local user, *\$home* is set from the password data, and can be tested in other preconditions that are evaluated after this one (the order of evaluation is given in section 3.12). However, the value of *\$home* can be overridden by **router\_home\_directory**. If the local part is not a local user, the router is skipped.

If you want to check that the local part is either the name of a local user or matches something else, you cannot combine **check\_local\_user** with a setting of **local\_parts**, because that specifies the logical *and* of the two conditions. However, you can use a *passwd* lookup in a **local\_parts** setting to achieve this. For example:

```
local_parts = passwd;$local_part : lsearch;/etc/other/users
```

Note, however, that the side effects of **check\_local\_user** (such as setting up a home directory) do not occur when a *passwd* lookup is used in a **local\_parts** (or any other) precondition.

<b>condition</b>	Use: <i>routers</i> ‡	Type: <i>string</i> †	Default: <i>unset</i>
------------------	-----------------------	-----------------------	-----------------------

This option specifies a general precondition test that has to succeed for the router to be called. The **condition** option is the last precondition to be evaluated (see section 3.12). The string is expanded, and if the result is a forced failure, or an empty string, or one of the strings “0” or “no” or “false” (checked without regard to the case of the letters), the router is skipped, and the address is offered to the next one.

If the result is any other value, the router is run (as this is the last precondition to be evaluated, all the other preconditions must be true).

This option is unusual in that multiple **condition** options may be present. All **condition** options must succeed.

The **condition** option provides a means of applying custom conditions to the running of routers. Note that in the case of a simple conditional expansion, the default expansion values are exactly what is wanted. For example:

```
condition = ${if >${message_age}{600}}
```

Because of the default behaviour of the string expansion, this is equivalent to

```
condition = ${if >${message_age}{600}{true}{}}
```

A multiple condition example, which succeeds:

```
condition = ${if >${message_age}{600}}
condition = ${if !eq${lc:$local_part}}{postmaster}}
condition = foobar
```

If the expansion fails (other than forced failure) delivery is deferred. Some of the other precondition options are common special cases that could in fact be specified using **condition**.

Historical note: We have **condition** on ACLs and on Routers. Routers are far older, and use one set of semantics. ACLs are newer and when they were created, the ACL **condition** process was given far stricter parse semantics. The **bool{}** expansion condition uses the same rules as ACLs. The **bool\_lax{}** expansion condition uses the same rules as Routers. More pointedly, the **bool\_lax{}** was written to match the existing Router rules processing behavior.

This is best illustrated in an example:

```
# If used in an ACL condition will fail with a syntax error, but
# in a router condition any extra characters are treated as a string

$ exim -be '${if eq ${lc:GOOGLE.com}} {google.com}} {yes} {no}}'
true {yes} {no}}

$ exim -be '${if eq ${lc:WHOIS.com}} {google.com}} {yes} {no}}'
{yes} {no}}
```

In each example above, the **if** statement actually ends after “{google.com}”. Since no true or false braces were defined, the default **if** behavior is to return a boolean true or a null answer (which evaluates to false). The rest of the line is then treated as a string. So the first example resulted in the boolean answer “true” with the string “ {yes} {no}” appended to it. The second example resulted in the null output (indicating false) with the string “ {yes} {no}” appended to it.

In fact you can put excess forward braces in too. In the router **condition**, Exim’s parser only looks for “{” symbols when they mean something, like after a “\$” or when required as part of a conditional. But otherwise “{” and “}” are treated as ordinary string characters.

Thus, in a Router, the above expansion strings will both always evaluate true, as the result of expansion is a non-empty string which doesn’t match an explicit false value. This can be tricky to debug. By contrast, in an ACL either of those strings will always result in an expansion error because the result doesn’t look sufficiently boolean.

<b>debug_print</b>	Use: <i>routers</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
--------------------	---------------------	----------------------------------	-----------------------

If this option is set and debugging is enabled (see the **-d** command line option) or in address-testing mode (see the **-bt** command line option), the string is expanded and included in the debugging output. If expansion of the string fails, the error message is written to the debugging output, and Exim carries on processing. This option is provided to help with checking out the values of variables and so on when debugging router configurations. For example, if a **condition** option appears not to be working, **debug\_print** can be used to output the variables it references. The output happens after checks for **domains**, **local\_parts**, and **check\_local\_user** but before any other preconditions are tested. A

newline is added to the text if it does not end with one. The variable *\$router\_name* contains the name of the router.

<b>disable_logging</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	---------------------	----------------------	-----------------------

If this option is set true, nothing is logged for any routing errors or for any deliveries caused by this router. You should not set this option unless you really, really know what you are doing. See also the generic transport option of the same name.

<b>dnssec_request_domains</b>	Use: <i>routers</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
-------------------------------	---------------------	---------------------------	-----------------------

DNS lookups for domains matching **dnssec\_request\_domains** will be done with the dnssec request bit set. This applies to all of the SRV, MX, AAAA, A lookup sequence.

<b>dnssec_require_domains</b>	Use: <i>routers</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
-------------------------------	---------------------	---------------------------	-----------------------

DNS lookups for domains matching **dnssec\_require\_domains** will be done with the dnssec request bit set. Any returns not having the Authenticated Data bit (AD bit) set will be ignored and logged as a host-lookup failure. This applies to all of the SRV, MX, AAAA, A lookup sequence.

<b>domains</b>	Use: <i>routers‡</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
----------------	----------------------	---------------------------	-----------------------

If this option is set, the router is skipped unless the current domain matches the list. If the match is achieved by means of a file lookup, the data that the lookup returned for the domain is placed in *\$domain\_data* for use in string expansions of the driver's private options. See section 3.12 for a list of the order in which preconditions are evaluated.

<b>driver</b>	Use: <i>routers</i>	Type: <i>string</i>	Default: <i>unset</i>
---------------	---------------------	---------------------	-----------------------

This option must always be set. It specifies which of the available routers is to be used.

<b>dsn_lasthop</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	---------------------	----------------------	-----------------------

If this option is set true, and extended DSN (RFC3461) processing is in effect, Exim will not pass on DSN requests to downstream DSN-aware hosts but will instead send a success DSN as if the next hop does not support DSN. Not effective on redirect routers.

<b>errors_to</b>	Use: <i>routers</i>	Type: <i>string†</i>	Default: <i>unset</i>
------------------	---------------------	----------------------	-----------------------

If a router successfully handles an address, it may assign the address to a transport for delivery or it may generate child addresses. In both cases, if there is a delivery problem during later processing, the resulting bounce message is sent to the address that results from expanding this string, provided that the address verifies successfully. The **errors\_to** option is expanded before **headers\_add**, **headers\_remove**, and **transport**.

The **errors\_to** setting associated with an address can be overridden if it subsequently passes through other routers that have their own **errors\_to** settings, or if the message is delivered by a transport with a **return\_path** setting.

If **errors\_to** is unset, or the expansion is forced to fail, or the result of the expansion fails to verify, the errors address associated with the incoming address is used. At top level, this is the envelope sender. A non-forced expansion failure causes delivery to be deferred.

If an address for which **errors\_to** has been set ends up being delivered over SMTP, the envelope sender for that delivery is the **errors\_to** value, so that any bounces that are generated by other MTAs



on the delivery route are also sent there. You can set **errors\_to** to the empty string by either of these settings:

```
errors_to =  
errors_to = ""
```

An expansion item that yields an empty string has the same effect. If you do this, a locally detected delivery error for addresses processed by this router no longer gives rise to a bounce message; the error is discarded. If the address is delivered to a remote host, the return path is set to <>, unless overridden by the **return\_path** option on the transport.

If for some reason you want to discard local errors, but use a non-empty MAIL command for remote delivery, you can preserve the original return path in *\$address\_data* in the router, and reinstate it in the transport by setting **return\_path**.

The most common use of **errors\_to** is to direct mailing list bounces to the manager of the list, as described in section 50.2, or to implement VERP (Variable Envelope Return Paths) (see section 50.6).

<b>expn</b>	Use: <i>routers</i> ‡	Type: <i>boolean</i>	Default: <i>true</i>
-------------	-----------------------	----------------------	----------------------

If this option is turned off, the router is skipped when testing an address as a result of processing an SMTP EXPN command. You might, for example, want to turn it off on a router for users' *.forward* files, while leaving it on for the system alias file. See section 3.12 for a list of the order in which preconditions are evaluated.

The use of the SMTP EXPN command is controlled by an ACL (see chapter 43). When Exim is running an EXPN command, it is similar to testing an address with **-bt**. Compare VRFY, whose counterpart is **-bv**.

<b>fail_verify</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	---------------------	----------------------	-----------------------

Setting this option has the effect of setting both **fail\_verify\_sender** and **fail\_verify\_recipient** to the same value.

<b>fail_verify_recipient</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------------	---------------------	----------------------	-----------------------

If this option is true and an address is accepted by this router when verifying a recipient, verification fails.

<b>fail_verify_sender</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	---------------------	----------------------	-----------------------

If this option is true and an address is accepted by this router when verifying a sender, verification fails.

<b>fallback_hosts</b>	Use: <i>routers</i>	Type: <i>string list</i>	Default: <i>unset</i>
-----------------------	---------------------	--------------------------	-----------------------

String expansion is not applied to this option. The argument must be a colon-separated list of host names or IP addresses. The list separator can be changed (see section 6.20), and a port can be specified with each name or address. In fact, the format of each item is exactly the same as defined for the list of hosts in a *manualroute* router (see section 20.5).

If a router queues an address for a remote transport, this host list is associated with the address, and used instead of the transport's fallback host list. If **hosts\_randomize** is set on the transport, the order of the list is randomized for each use. See the **fallback\_hosts** option of the *smtp* transport for further details.

<b>group</b>	Use: <i>routers</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
--------------	---------------------	----------------------------------	---------------------------

When a router queues an address for a transport, and the transport does not specify a group, the group given here is used when running the delivery process. The group may be specified numerically or by name. If expansion fails, the error is logged and delivery is deferred. The default is unset, unless **check\_local\_user** is set, when the default is taken from the password information. See also **initgroups** and **user** and the discussion in chapter 23.

<b>headers_add</b>	Use: <i>routers</i>	Type: <i>list</i> <sup>†</sup>	Default: <i>unset</i>
--------------------	---------------------	--------------------------------	-----------------------

This option specifies a list of text headers, newline-separated (by default, changeable in the usual way), that is associated with any addresses that are accepted by the router. Each item is separately expanded, at routing time. However, this option has no effect when an address is just being verified. The way in which the text is used to add header lines at transport time is described in section 47.17. New header lines are not actually added until the message is in the process of being transported. This means that references to header lines in string expansions in the transport's configuration do not "see" the added header lines.

The **headers\_add** option is expanded after **errors\_to**, but before **headers\_remove** and **transport**. If an item is empty, or if an item expansion is forced to fail, the item has no effect. Other expansion failures are treated as configuration errors.

Unlike most options, **headers\_add** can be specified multiple times for a router; all listed headers are added.

**Warning 1:** The **headers\_add** option cannot be used for a *redirect* router that has the **one\_time** option set.

**Warning 2:** If the **unseen** option is set on the router, all header additions are deleted when the address is passed on to subsequent routers. For a **redirect** router, if a generated address is the same as the incoming address, this can lead to duplicate addresses with different header modifications. Exim does not do duplicate deliveries (except, in certain circumstances, to pipes -- see section 22.7), but it is undefined which of the duplicates is discarded, so this ambiguous situation should be avoided. The **repeat\_use** option of the **redirect** router may be of help.

<b>headers_remove</b>	Use: <i>routers</i>	Type: <i>list</i> <sup>†</sup>	Default: <i>unset</i>
-----------------------	---------------------	--------------------------------	-----------------------

This option specifies a list of text headers, colon-separated (by default, changeable in the usual way), that is associated with any addresses that are accepted by the router. Each item is separately expanded, at routing time. However, this option has no effect when an address is just being verified. The way in which the text is used to remove header lines at transport time is described in section 47.17. Header lines are not actually removed until the message is in the process of being transported. This means that references to header lines in string expansions in the transport's configuration still "see" the original header lines.

The **headers\_remove** option is expanded after **errors\_to** and **headers\_add**, but before **transport**. If an item expansion is forced to fail, the item has no effect. Other expansion failures are treated as configuration errors.

Unlike most options, **headers\_remove** can be specified multiple times for a router; all listed headers are removed.

**Warning 1:** The **headers\_remove** option cannot be used for a *redirect* router that has the **one\_time** option set.

**Warning 2:** If the **unseen** option is set on the router, all header removal requests are deleted when the address is passed on to subsequent routers, and this can lead to problems with duplicates -- see the similar warning for **headers\_add** above.

**Warning 3:** Because of the separate expansion of the list items, items that contain a list separator must have it doubled. To avoid this, change the list separator (6.21).

<b>ignore_target_hosts</b>	Use: <i>routers</i>	Type: <i>host list</i> †	Default: <i>unset</i>
----------------------------	---------------------	--------------------------	-----------------------

Although this option is a host list, it should normally contain IP address entries rather than names. If any host that is looked up by the router has an IP address that matches an item in this list, Exim behaves as if that IP address did not exist. This option allows you to cope with rogue DNS entries like

```
remote.domain.example.  A  127.0.0.1
```

by setting

```
ignore_target_hosts = 127.0.0.1
```

on the relevant router. If all the hosts found by a *dnslookup* router are discarded in this way, the router declines. In a conventional configuration, an attempt to mail to such a domain would normally provoke the “unrouteable domain” error, and an attempt to verify an address in the domain would fail. Similarly, if **ignore\_target\_hosts** is set on an *ipliteral* router, the router declines if presented with one of the listed addresses.

You can use this option to disable the use of IPv4 or IPv6 for mail delivery by means of the first or the second of the following settings, respectively:

```
ignore_target_hosts = 0.0.0.0/0
ignore_target_hosts = <; 0::0/0
```

The pattern in the first line matches all IPv4 addresses, whereas the pattern in the second line matches all IPv6 addresses.

This option may also be useful for ignoring link-local and site-local IPv6 addresses. Because, like all host lists, the value of **ignore\_target\_hosts** is expanded before use as a list, it is possible to make it dependent on the domain that is being routed.

During its expansion, *\$host\_address* is set to the IP address that is being checked.

<b>initgroups</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------	---------------------	----------------------	-----------------------

If the router queues an address for a transport, and this option is true, and the uid supplied by the router is not overridden by the transport, the *initgroups()* function is called when running the transport to ensure that any additional groups associated with the uid are set up. See also **group** and **user** and the discussion in chapter 23.

<b>local_part_prefix</b>	Use: <i>routers</i> ‡	Type: <i>string list</i>	Default: <i>unset</i>
--------------------------	-----------------------	--------------------------	-----------------------

If this option is set, the router is skipped unless the local part starts with one of the given strings, or **local\_part\_prefix\_optional** is true. See section 3.12 for a list of the order in which preconditions are evaluated.

The list is scanned from left to right, and the first prefix that matches is used. A limited form of wildcard is available; if the prefix begins with an asterisk, it matches the longest possible sequence of arbitrary characters at the start of the local part. An asterisk should therefore always be followed by some character that does not occur in normal local parts. Wildcarding can be used to set up multiple user mailboxes, as described in section 50.8.

During the testing of the **local\_parts** option, and while the router is running, the prefix is removed from the local part, and is available in the expansion variable *\$local\_part\_prefix*. When a message is being delivered, if the router accepts the address, this remains true during subsequent delivery by a transport. In particular, the local part that is transmitted in the RCPT command for LMTP, SMTP, and BSMTP deliveries has the prefix removed by default. This behaviour can be overridden by setting **rcpt\_include\_affixes** true on the relevant transport.

When an address is being verified, **local\_part\_prefix** affects only the behaviour of the router. If the callout feature of verification is in use, this means that the full address, including the prefix, will be used during the callout.

The prefix facility is commonly used to handle local parts of the form **owner-something**. Another common use is to support local parts of the form **real-username** to bypass a user's *.forward* file – helpful when trying to tell a user their forwarding is broken – by placing a router like this one immediately before the router that handles *.forward* files:

```
real_localuser:
    driver = accept
    local_part_prefix = real-
    check_local_user
    transport = local_delivery
```

For security, it would probably be a good idea to restrict the use of this router to locally-generated messages, using a condition such as this:

```
condition = ${if match {$sender_host_address}\
                    {\N^(|127\.0\.0\.1)\N$}}
```

If both **local\_part\_prefix** and **local\_part\_suffix** are set for a router, both conditions must be met if not optional. Care must be taken if wildcards are used in both a prefix and a suffix on the same router. Different separator characters must be used to avoid ambiguity.

<b>local_part_prefix_optional</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------------------	---------------------	----------------------	-----------------------

See **local\_part\_prefix** above.

<b>local_part_suffix</b>	Use: <i>routers‡</i>	Type: <i>string list</i>	Default: <i>unset</i>
--------------------------	----------------------	--------------------------	-----------------------

This option operates in the same way as **local\_part\_prefix**, except that the local part must end (rather than start) with the given string, the **local\_part\_suffix\_optional** option determines whether the suffix is mandatory, and the wildcard *\** character, if present, must be the last character of the suffix. This option facility is commonly used to handle local parts of the form **something-request** and multiple user mailboxes of the form **username-foo**.

<b>local_part_suffix_optional</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------------------	---------------------	----------------------	-----------------------

See **local\_part\_suffix** above.

<b>local_parts</b>	Use: <i>routers‡</i>	Type: <i>local part list†</i>	Default: <i>unset</i>
--------------------	----------------------	-------------------------------	-----------------------

The router is run only if the local part of the address matches the list. See section 3.12 for a list of the order in which preconditions are evaluated, and section 10.21 for a discussion of local part lists. Because the string is expanded, it is possible to make it depend on the domain, for example:

```
local_parts = dbm;/usr/local/specials/$domain
```

If the match is achieved by a lookup, the data that the lookup returned for the local part is placed in the variable *\$local\_part\_data* for use in expansions of the router's private options. You might use this option, for example, if you have a large number of local virtual domains, and you want to send all postmaster mail to the same place without having to set up an alias in each virtual domain:

```
postmaster:
    driver = redirect
    local_parts = postmaster
    data = postmaster@real.domain.example
```

<b>log_as_local</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>see below</i>
---------------------	---------------------	----------------------	---------------------------

Exim has two logging styles for delivery, the idea being to make local deliveries stand out more visibly from remote ones. In the “local” style, the recipient address is given just as the local part, without a domain. The use of this style is controlled by this option. It defaults to true for the *accept* router, and false for all the others. This option applies only when a router assigns an address to a transport. It has no effect on routers that redirect addresses.

<b>more</b>	Use: <i>routers</i>	Type: <i>boolean†</i>	Default: <i>true</i>
-------------	---------------------	-----------------------	----------------------

The result of string expansion for this option must be a valid boolean value, that is, one of the strings “yes”, “no”, “true”, or “false”. Any other result causes an error, and delivery is deferred. If the expansion is forced to fail, the default value for the option (true) is used. Other failures cause delivery to be deferred.

If this option is set false, and the router declines to handle the address, no further routers are tried, routing fails, and the address is bounced. However, if the router explicitly passes an address to the following router by means of the setting

```
self = pass
```

or otherwise, the setting of **more** is ignored. Also, the setting of **more** does not affect the behaviour if one of the precondition tests fails. In that case, the address is always passed to the next router.

Note that **address\_data** is not considered to be a precondition. If its expansion is forced to fail, the router declines, and the value of **more** controls what happens next.

<b>pass_on_timeout</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	---------------------	----------------------	-----------------------

If a router times out during a host lookup, it normally causes deferral of the address. If **pass\_on\_timeout** is set, the address is passed on to the next router, overriding **no\_more**. This may be helpful for systems that are intermittently connected to the Internet, or those that want to pass to a smart host any messages that cannot immediately be delivered.

There are occasional other temporary errors that can occur while doing DNS lookups. They are treated in the same way as a timeout, and this option applies to all of them.

<b>pass_router</b>	Use: <i>routers</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------------	---------------------	---------------------	-----------------------

Routers that recognize the generic **self** option (*dnslookup*, *ipliteral*, and *manualroute*) are able to return “pass”, forcing routing to continue, and overriding a false setting of **more**. When one of these routers returns “pass”, the address is normally handed on to the next router in sequence. This can be changed by setting **pass\_router** to the name of another router. However (unlike **redirect\_router**) the named router must be below the current router, to avoid loops. Note that this option applies only to the special case of “pass”. It does not apply when a router returns “decline” because it cannot handle an address.

<b>redirect_router</b>	Use: <i>routers</i>	Type: <i>string</i>	Default: <i>unset</i>
------------------------	---------------------	---------------------	-----------------------

Sometimes an administrator knows that it is pointless to reprocess addresses generated from alias or forward files with the same router again. For example, if an alias file translates real names into login ids there is no point searching the alias file a second time, especially if it is a large file.

The **redirect\_router** option can be set to the name of any router instance. It causes the routing of any generated addresses to start at the named router instead of at the first router. This option has no effect if the router in which it is set does not generate new addresses.

<b>require_files</b>	Use: <i>routers</i> ‡	Type: <i>string list</i> †	Default: <i>unset</i>
----------------------	-----------------------	----------------------------	-----------------------

This option provides a general mechanism for predicating the running of a router on the existence or non-existence of certain files or directories. Before running a router, as one of its precondition tests, Exim works its way through the **require\_files** list, expanding each item separately.

Because the list is split before expansion, any colons in expansion items must be doubled, or the facility for using a different list separator must be used. If any expansion is forced to fail, the item is ignored. Other expansion failures cause routing of the address to be deferred.

If any expanded string is empty, it is ignored. Otherwise, except as described below, each string must be a fully qualified file path, optionally preceded by “!”. The paths are passed to the *stat()* function to test for the existence of the files or directories. The router is skipped if any paths not preceded by “!” do not exist, or if any paths preceded by “!” do exist.

If *stat()* cannot determine whether a file exists or not, delivery of the message is deferred. This can happen when NFS-mounted filesystems are unavailable.

This option is checked after the **domains**, **local\_parts**, and **senders** options, so you cannot use it to check for the existence of a file in which to look up a domain, local part, or sender. (See section 3.12 for a full list of the order in which preconditions are evaluated.) However, as these options are all expanded, you can use the **exists** expansion condition to make such tests. The **require\_files** option is intended for checking files that the router may be going to use internally, or which are needed by a transport (for example *.procmailrc*).

During delivery, the *stat()* function is run as root, but there is a facility for some checking of the accessibility of a file by another user. This is not a proper permissions check, but just a “rough” check that operates as follows:

If an item in a **require\_files** list does not contain any forward slash characters, it is taken to be the user (and optional group, separated by a comma) to be checked for subsequent files in the list. If no group is specified but the user is specified symbolically, the gid associated with the uid is used. For example:

```
require_files = mail:/some/file
require_files = $local_part:$home/.procmailrc
```

If a user or group name in a **require\_files** list does not exist, the **require\_files** condition fails.

Exim performs the check by scanning along the components of the file path, and checking the access for the given uid and gid. It checks for “x” access on directories, and “r” access on the final file. Note that this means that file access control lists, if the operating system has them, are ignored.

**Warning 1:** When the router is being run to verify addresses for an incoming SMTP message, Exim is not running as root, but under its own uid. This may affect the result of a **require\_files** check. In particular, *stat()* may yield the error EACCES (“Permission denied”). This means that the Exim user is not permitted to read one of the directories on the file’s path.

**Warning 2:** Even when Exim is running as root while delivering a message, *stat()* can yield EACCES for a file in an NFS directory that is mounted without root access. In this case, if a check for access by a particular user is requested, Exim creates a subprocess that runs as that user, and tries the check again in that process.

The default action for handling an unresolved EACCES is to consider it to be caused by a configuration error, and routing is deferred because the existence or non-existence of the file cannot be determined. However, in some circumstances it may be desirable to treat this condition as if the file did not exist. If the file name (or the exclamation mark that precedes the file name for non-existence) is preceded by a plus sign, the EACCES error is treated as if the file did not exist. For example:

```
require_files = +/some/file
```

If the router is not an essential part of verification (for example, it handles users’ *.forward* files), another solution is to set the **verify** option false so that the router is skipped when verifying.

<b>retry_use_local_part</b>	Use: <i>routers</i>	Type: <i>boolean</i>	Default: <i>see below</i>
-----------------------------	---------------------	----------------------	---------------------------

When a delivery suffers a temporary routing failure, a retry record is created in Exim's hints database. For addresses whose routing depends only on the domain, the key for the retry record should not involve the local part, but for other addresses, both the domain and the local part should be included. Usually, remote routing is of the former kind, and local routing is of the latter kind.

This option controls whether the local part is used to form the key for retry hints for addresses that suffer temporary errors while being handled by this router. The default value is true for any router that has **check\_local\_user** set, and false otherwise. Note that this option does not apply to hints keys for transport delays; they are controlled by a generic transport option of the same name.

The setting of **retry\_use\_local\_part** applies only to the router on which it appears. If the router generates child addresses, they are routed independently; this setting does not become attached to them.

<b>router_home_directory</b>	Use: <i>routers</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
------------------------------	---------------------	----------------------------------	-----------------------

This option sets a home directory for use while the router is running. (Compare **transport\_home\_directory**, which sets a home directory for later transporting.) In particular, if used on a *redirect* router, this option sets a value for *\$home* while a filter is running. The value is expanded; forced expansion failure causes the option to be ignored – other failures cause the router to defer.

Expansion of **router\_home\_directory** happens immediately after the **check\_local\_user** test (if configured), before any further expansions take place. (See section 3.12 for a list of the order in which preconditions are evaluated.) While the router is running, **router\_home\_directory** overrides the value of *\$home* that came from **check\_local\_user**.

When a router accepts an address and assigns it to a local transport (including the cases when a *redirect* router generates a pipe, file, or autoreply delivery), the home directory setting for the transport is taken from the first of these values that is set:

- The **home\_directory** option on the transport;
- The **transport\_home\_directory** option on the router;
- The password data if **check\_local\_user** is set on the router;
- The **router\_home\_directory** option on the router.

In other words, **router\_home\_directory** overrides the password data for the router, but not for the transport.

<b>self</b>	Use: <i>routers</i>	Type: <i>string</i>	Default: <i>freeze</i>
-------------	---------------------	---------------------	------------------------

This option applies to those routers that use a recipient address to find a list of remote hosts. Currently, these are the *dnslookup*, *ipliteral*, and *manualroute* routers. Certain configurations of the *queryprogram* router can also specify a list of remote hosts. Usually such routers are configured to send the message to a remote host via an *smtp* transport. The **self** option specifies what happens when the first host on the list turns out to be the local host. The way in which Exim checks for the local host is described in section 13.8.

Normally this situation indicates either an error in Exim's configuration (for example, the router should be configured not to process this domain), or an error in the DNS (for example, the MX should not point to this host). For this reason, the default action is to log the incident, defer the address, and freeze the message. The following alternatives are provided for use in special cases:

#### **defer**

Delivery of the message is tried again later, but the message is not frozen.

**reroute:** <domain>

The domain is changed to the given domain, and the address is passed back to be reprocessed by the routers. No rewriting of headers takes place. This behaviour is essentially a redirection.

**reroute: rewrite:** <domain>

The domain is changed to the given domain, and the address is passed back to be reprocessed by the routers. Any headers that contain the original domain are rewritten.

**pass**

The router passes the address to the next router, or to the router named in the **pass\_router** option if it is set. This overrides **no\_more**. During subsequent routing and delivery, the variable *\$self\_hostname* contains the name of the local host that the router encountered. This can be used to distinguish between different cases for hosts with multiple names. The combination

```
self = pass
no_more
```

ensures that only those addresses that routed to the local host are passed on. Without **no\_more**, addresses that were declined for other reasons would also be passed to the next router.

**fail**

Delivery fails and an error report is generated.

**send**

The anomaly is ignored and the address is queued for the transport. This setting should be used with extreme caution. For an *smtp* transport, it makes sense only in cases where the program that is listening on the SMTP port is not this version of Exim. That is, it must be some other MTA, or Exim with a different configuration file that handles the domain in another way.

<b>senders</b>	Use: <i>routers</i> ‡	Type: <i>address list</i> †	Default: <i>unset</i>
----------------	-----------------------	-----------------------------	-----------------------

If this option is set, the router is skipped unless the message's sender address matches something on the list. See section 3.12 for a list of the order in which preconditions are evaluated.

There are issues concerning verification when the running of routers is dependent on the sender. When Exim is verifying the address in an **errors\_to** setting, it sets the sender to the null string. When using the **-bt** option to check a configuration file, it is necessary also to use the **-f** option to set an appropriate sender. For incoming mail, the sender is unset when verifying the sender, but is available when verifying any recipients. If the SMTP VRFY command is enabled, it must be used after MAIL if the sender address matters.

<b>translate_ip_address</b>	Use: <i>routers</i>	Type: <i>string</i> †	Default: <i>unset</i>
-----------------------------	---------------------	-----------------------	-----------------------

There exist some rare networking situations (for example, packet radio) where it is helpful to be able to translate IP addresses generated by normal routing mechanisms into other IP addresses, thus performing a kind of manual IP routing. This should be done only if the normal IP routing of the TCP/IP stack is inadequate or broken. Because this is an extremely uncommon requirement, the code to support this option is not included in the Exim binary unless `SUPPORT_TRANSLATE_IP_ADDRESS=yes` is set in *Local/Makefile*.

The **translate\_ip\_address** string is expanded for every IP address generated by the router, with the generated address set in *\$host\_address*. If the expansion is forced to fail, no action is taken. For any other expansion error, delivery of the message is deferred. If the result of the expansion is an IP address, that replaces the original address; otherwise the result is assumed to be a host name – this is looked up using *gethostbyname()* (or *getipnodebyname()* when available) to produce one or more replacement IP addresses. For example, to subvert all IP addresses in some specific networks, this could be added to a router:

```
translate_ip_address = \
  ${lookup{${mask:$host_address/26}}lsearch{/some/file}\
  {$value}fail}}
```



The file would contain lines like

```
10.2.3.128/26    some.host
10.8.4.34/26     10.44.8.15
```

You should not make use of this facility unless you really understand what you are doing.

<b>transport</b>	Use: <i>routers</i>	Type: <i>string†</i>	Default: <i>unset</i>
------------------	---------------------	----------------------	-----------------------

This option specifies the transport to be used when a router accepts an address and sets it up for delivery. A transport is never needed if a router is used only for verification. The value of the option is expanded at routing time, after the expansion of **errors\_to**, **headers\_add**, and **headers\_remove**, and result must be the name of one of the configured transports. If it is not, delivery is deferred.

The **transport** option is not used by the *redirect* router, but it does have some private options that set up transports for pipe and file deliveries (see chapter 22).

<b>transport_current_directory</b>	Use: <i>routers</i>	Type: <i>string†</i>	Default: <i>unset</i>
------------------------------------	---------------------	----------------------	-----------------------

This option associates a current directory with any address that is routed to a local transport. This can happen either because a transport is explicitly configured for the router, or because it generates a delivery to a file or a pipe. During the delivery process (that is, at transport time), this option string is expanded and is set as the current directory, unless overridden by a setting on the transport. If the expansion fails for any reason, including forced failure, an error is logged, and delivery is deferred. See chapter 23 for details of the local delivery environment.

<b>transport_home_directory</b>	Use: <i>routers</i>	Type: <i>string†</i>	Default: <i>see below</i>
---------------------------------	---------------------	----------------------	---------------------------

This option associates a home directory with any address that is routed to a local transport. This can happen either because a transport is explicitly configured for the router, or because it generates a delivery to a file or a pipe. During the delivery process (that is, at transport time), the option string is expanded and is set as the home directory, unless overridden by a setting of **home\_directory** on the transport. If the expansion fails for any reason, including forced failure, an error is logged, and delivery is deferred.

If the transport does not specify a home directory, and **transport\_home\_directory** is not set for the router, the home directory for the transport is taken from the password data if **check\_local\_user** is set for the router. Otherwise it is taken from **router\_home\_directory** if that option is set; if not, no home directory is set for the transport.

See chapter 23 for further details of the local delivery environment.

<b>unseen</b>	Use: <i>routers</i>	Type: <i>boolean†</i>	Default: <i>false</i>
---------------	---------------------	-----------------------	-----------------------

The result of string expansion for this option must be a valid boolean value, that is, one of the strings “yes”, “no”, “true”, or “false”. Any other result causes an error, and delivery is deferred. If the expansion is forced to fail, the default value for the option (false) is used. Other failures cause delivery to be deferred.

When this option is set true, routing does not cease if the router accepts the address. Instead, a copy of the incoming address is passed to the next router, overriding a false setting of **more**. There is little point in setting **more** false if **unseen** is always true, but it may be useful in cases when the value of **unseen** contains expansion items (and therefore, presumably, is sometimes true and sometimes false).

Setting the **unseen** option has a similar effect to the **unseen** command qualifier in filter files. It can be used to cause copies of messages to be delivered to some other destination, while also carrying out a normal delivery. In effect, the current address is made into a “parent” that has two children – one that is delivered as specified by this router, and a clone that goes on to be routed further. For this reason, **unseen** may not be combined with the **one\_time** option in a *redirect* router.

**Warning:** Header lines added to the address (or specified for removal) by this router or by previous routers affect the “unseen” copy of the message only. The clone that continues to be processed by further routers starts with no added headers and none specified for removal. For a **redirect** router, if a generated address is the same as the incoming address, this can lead to duplicate addresses with different header modifications. Exim does not do duplicate deliveries (except, in certain circumstances, to pipes -- see section 22.7), but it is undefined which of the duplicates is discarded, so this ambiguous situation should be avoided. The **repeat\_use** option of the **redirect** router may be of help.

Unlike the handling of header modifications, any data that was set by the **address\_data** option in the current or previous routers *is* passed on to subsequent routers.

<b>user</b>	Use: <i>routers</i>	Type: <i>string†</i>	Default: <i>see below</i>
-------------	---------------------	----------------------	---------------------------

When a router queues an address for a transport, and the transport does not specify a user, the user given here is used when running the delivery process. The user may be specified numerically or by name. If expansion fails, the error is logged and delivery is deferred. This user is also used by the *redirect* router when running a filter file. The default is unset, except when **check\_local\_user** is set. In this case, the default is taken from the password information. If the user is specified as a name, and **group** is not set, the group associated with the user is used. See also **initgroups** and **group** and the discussion in chapter 23.

<b>verify</b>	Use: <i>routers‡</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------	----------------------	----------------------	----------------------

Setting this option has the effect of setting **verify\_sender** and **verify\_recipient** to the same value.

<b>verify_only</b>	Use: <i>routers‡</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	----------------------	----------------------	-----------------------

If this option is set, the router is used only when verifying an address, delivering in cutthrough mode or testing with the **-bv** option, not when actually doing a delivery, testing with the **-bt** option, or running the SMTP EXPN command. It can be further restricted to verifying only senders or recipients by means of **verify\_sender** and **verify\_recipient**.

**Warning:** When the router is being run to verify addresses for an incoming SMTP message, Exim is not running as root, but under its own uid. If the router accesses any files, you need to make sure that they are accessible to the Exim user or group.

<b>verify_recipient</b>	Use: <i>routers‡</i>	Type: <i>boolean</i>	Default: <i>true</i>
-------------------------	----------------------	----------------------	----------------------

If this option is false, the router is skipped when verifying recipient addresses, delivering in cutthrough mode or testing recipient verification using **-bv**. See section 3.12 for a list of the order in which preconditions are evaluated. See also the *\$verify\_mode* variable.

<b>verify_sender</b>	Use: <i>routers‡</i>	Type: <i>boolean</i>	Default: <i>true</i>
----------------------	----------------------	----------------------	----------------------

If this option is false, the router is skipped when verifying sender addresses or testing sender verification using **-bvs**. See section 3.12 for a list of the order in which preconditions are evaluated. See also the *\$verify\_mode* variable.

## 16. The accept router

The *accept* router has no private options of its own. Unless it is being used purely for verification (see **verify\_only**) a transport is required to be defined by the generic **transport** option. If the preconditions that are specified by generic options are met, the router accepts the address and queues it for the given transport. The most common use of this router is for setting up deliveries to local mailboxes. For example:

```
localusers:
  driver = accept
  domains = mydomain.example
  check_local_user
  transport = local_delivery
```

The **domains** condition in this example checks the domain of the address, and **check\_local\_user** checks that the local part is the login of a local user. When both preconditions are met, the *accept* router runs, and queues the address for the *local\_delivery* transport.

## 17. The *dnslookup* router

The *dnslookup* router looks up the hosts that handle mail for the recipient's domain in the DNS. A transport must always be set for this router, unless **verify\_only** is set.

If SRV support is configured (see **check\_srv** below), Exim first searches for SRV records. If none are found, or if SRV support is not configured, MX records are looked up. If no MX records exist, address records are sought. However, **mx\_domains** can be set to disable the direct use of address records.

MX records of equal priority are sorted by Exim into a random order. Exim then looks for address records for the host names obtained from MX or SRV records. When a host has more than one IP address, they are sorted into a random order, except that IPv6 addresses are always sorted before IPv4 addresses. If all the IP addresses found are discarded by a setting of the **ignore\_target\_hosts** generic option, the router declines.

Unless they have the highest priority (lowest MX value), MX records that point to the local host, or to any host name that matches **hosts\_treat\_as\_local**, are discarded, together with any other MX records of equal or lower priority.

If the host pointed to by the highest priority MX record, or looked up as an address record, is the local host, or matches **hosts\_treat\_as\_local**, what happens is controlled by the generic **self** option.

### 17.1 Problems with DNS lookups

There have been problems with DNS servers when SRV records are looked up. Some misbehaving servers return a DNS error or timeout when a non-existent SRV record is sought. Similar problems have in the past been reported for MX records. The global **dns\_again\_means\_nonexist** option can help with this problem, but it is heavy-handed because it is a global option.

For this reason, there are two options, **srv\_fail\_domains** and **mx\_fail\_domains**, that control what happens when a DNS lookup in a *dnslookup* router results in a DNS failure or a "try again" response. If an attempt to look up an SRV or MX record causes one of these results, and the domain matches the relevant list, Exim behaves as if the DNS had responded "no such record". In the case of an SRV lookup, this means that the router proceeds to look for MX records; in the case of an MX lookup, it proceeds to look for A or AAAA records, unless the domain matches **mx\_domains**, in which case routing fails.

### 17.2 Declining addresses by *dnslookup*

There are a few cases where a *dnslookup* router will decline to accept an address; if such a router is expected to handle "all remaining non-local domains", then it is important to set **no\_more**.

The router will defer rather than decline if the domain is found in the **fail\_defer\_domains** router option.

Reasons for a *dnslookup* router to decline currently include:

- The domain does not exist in DNS
- The domain exists but the MX record's host part is just "."; this is a common convention (borrowed from SRV) used to indicate that there is no such service for this domain and to not fall back to trying A/AAAA records.
- Ditto, but for SRV records, when **check\_srv** is set on this router.
- MX record points to a non-existent host.
- MX record points to an IP address and the main section option **allow\_mx\_to\_ip** is not set.
- MX records exist and point to valid hosts, but all hosts resolve only to addresses blocked by the **ignore\_target\_hosts** generic option on this router.

- The domain is not syntactically valid (see also **allow\_utf8\_domains** and **dns\_check\_names\_pattern** for handling one variant of this)
- **check\_secondary\_mx** is set on this router but the local host can not be found in the MX records (see below)

## 17.3 Private options for dnslookup

The private options for the *dnslookup* router are as follows:

<b>check_secondary_mx</b>	Use: <i>dnslookup</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	-----------------------	----------------------	-----------------------

If this option is set, the router declines unless the local host is found in (and removed from) the list of hosts obtained by MX lookup. This can be used to process domains for which the local host is a secondary mail exchanger differently to other domains. The way in which Exim decides whether a host is the local host is described in section 13.8.

<b>check_srv</b>	Use: <i>dnslookup</i>	Type: <i>string†</i>	Default: <i>unset</i>
------------------	-----------------------	----------------------	-----------------------

The *dnslookup* router supports the use of SRV records (see RFC 2782) in addition to MX and address records. The support is disabled by default. To enable SRV support, set the **check\_srv** option to the name of the service required. For example,

```
check_srv = smtp
```

looks for SRV records that refer to the normal smtp service. The option is expanded, so the service name can vary from message to message or address to address. This might be helpful if SRV records are being used for a submission service. If the expansion is forced to fail, the **check\_srv** option is ignored, and the router proceeds to look for MX records in the normal way.

When the expansion succeeds, the router searches first for SRV records for the given service (it assumes TCP protocol). A single SRV record with a host name that consists of just a single dot indicates “no such service for this domain”; if this is encountered, the router declines. If other kinds of SRV record are found, they are used to construct a host list for delivery according to the rules of RFC 2782. MX records are not sought in this case.

When no SRV records are found, MX records (and address records) are sought in the traditional way. In other words, SRV records take precedence over MX records, just as MX records take precedence over address records. Note that this behaviour is not sanctioned by RFC 2782, though a previous draft RFC defined it. It is apparently believed that MX records are sufficient for email and that SRV records should not be used for this purpose. However, SRV records have an additional “weight” feature which some people might find useful when trying to split an SMTP load between hosts of different power.

See section 17.1 above for a discussion of Exim’s behaviour when there is a DNS lookup error.

<b>fail_defer_domains</b>	Use: <i>dnslookup</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
---------------------------	-----------------------	---------------------------	-----------------------

DNS lookups for domains matching **fail\_defer\_domains** which find no matching record will cause the router to defer rather than the default behaviour of decline. This maybe be useful for queueing messages for a newly created domain while the DNS configuration is not ready. However, it will result in any message with mistyped domains also being queued.

<b>mx_domains</b>	Use: <i>dnslookup</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
-------------------	-----------------------	---------------------------	-----------------------

A domain that matches **mx\_domains** is required to have either an MX or an SRV record in order to be recognized. (The name of this option could be improved.) For example, if all the mail hosts in *fict.example* are known to have MX records, except for those in *discworld.fict.example*, you could use this setting:

```
mx_domains = ! *.discworld.fict.example : *.fict.example
```

This specifies that messages addressed to a domain that matches the list but has no MX record should be bounced immediately instead of being routed using the address record.

<b>mx_fail_domains</b>	Use: <i>dnslookup</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
------------------------	-----------------------	---------------------------	-----------------------

If the DNS lookup for MX records for one of the domains in this list causes a DNS lookup error, Exim behaves as if no MX records were found. See section 17.1 for more discussion.

<b>qualify_single</b>	Use: <i>dnslookup</i>	Type: <i>boolean</i>	Default: <i>true</i>
-----------------------	-----------------------	----------------------	----------------------

When this option is true, the resolver option RES\_DEFNAMES is set for DNS lookups. Typically, but not standardly, this causes the resolver to qualify single-component names with the default domain. For example, on a machine called *dictionary.ref.example*, the domain *thesaurus* would be changed to *thesaurus.ref.example* inside the resolver. For details of what your resolver actually does, consult your man pages for *resolver* and *resolv.conf*.

<b>rewrite_headers</b>	Use: <i>dnslookup</i>	Type: <i>boolean</i>	Default: <i>true</i>
------------------------	-----------------------	----------------------	----------------------

If the domain name in the address that is being processed is not fully qualified, it may be expanded to its full form by a DNS lookup. For example, if an address is specified as *dormouse@teaparty*, the domain might be expanded to *teaparty.wonderland.fict.example*. Domain expansion can also occur as a result of setting the **widen\_domains** option. If **rewrite\_headers** is true, all occurrences of the abbreviated domain name in any *Bcc:*, *Cc:*, *From:*, *Reply-to:*, *Sender:*, and *To:* header lines of the message are rewritten with the full domain name.

This option should be turned off only when it is known that no message is ever going to be sent outside an environment where the abbreviation makes sense.

When an MX record is looked up in the DNS and matches a wildcard record, name servers normally return a record containing the name that has been looked up, making it impossible to detect whether a wildcard was present or not. However, some name servers have recently been seen to return the wildcard entry. If the name returned by a DNS lookup begins with an asterisk, it is not used for header rewriting.

<b>same_domain_copy_routing</b>	Use: <i>dnslookup</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------------	-----------------------	----------------------	-----------------------

Addresses with the same domain are normally routed by the *dnslookup* router to the same list of hosts. However, this cannot be presumed, because the router options and preconditions may refer to the local part of the address. By default, therefore, Exim routes each address in a message independently. DNS servers run caches, so repeated DNS lookups are not normally expensive, and in any case, personal messages rarely have more than a few recipients.

If you are running mailing lists with large numbers of subscribers at the same domain, and you are using a *dnslookup* router which is independent of the local part, you can set **same\_domain\_copy\_routing** to bypass repeated DNS lookups for identical domains in one message. In this case, when *dnslookup* routes an address to a remote transport, any other unrouted addresses in the message that have the same domain are automatically given the same routing without processing them independently, provided the following conditions are met:

- No router that processed the address specified **headers\_add** or **headers\_remove**.
- The router did not change the address in any way, for example, by “widening” the domain.

<b>search_parents</b>	Use: <i>dnslookup</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	-----------------------	----------------------	-----------------------

When this option is true, the resolver option RES\_DNSRCH is set for DNS lookups. This is different from the **qualify\_single** option in that it applies to domains containing dots. Typically, but not standardly, it causes the resolver to search for the name in the current domain and in parent domains. For example, on a machine in the *fict.example* domain, if looking up *teaparty.wonderland* failed, the resolver would try *teaparty.wonderland.fict.example*. For details of what your resolver actually does, consult your man pages for *resolver* and *resolv.conf*.

Setting this option true can cause problems in domains that have a wildcard MX record, because any domain that does not have its own MX record matches the local wildcard.

<b>srv_fail_domains</b>	Use: <i>dnslookup</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
-------------------------	-----------------------	---------------------------	-----------------------

If the DNS lookup for SRV records for one of the domains in this list causes a DNS lookup error, Exim behaves as if no SRV records were found. See section 17.1 for more discussion.

<b>widen_domains</b>	Use: <i>dnslookup</i>	Type: <i>string list</i>	Default: <i>unset</i>
----------------------	-----------------------	--------------------------	-----------------------

If a DNS lookup fails and this option is set, each of its strings in turn is added onto the end of the domain, and the lookup is tried again. For example, if

```
widen_domains = fict.example:ref.example
```

is set and a lookup of *klington.dictionary* fails, *klington.dictionary.fict.example* is looked up, and if this fails, *klington.dictionary.ref.example* is tried. Note that the **qualify\_single** and **search\_parents** options can cause some widening to be undertaken inside the DNS resolver. **widen\_domains** is not applied to sender addresses when verifying, unless **rewrite\_headers** is false (not the default).

## 17.4 Effect of **qualify\_single** and **search\_parents**

When a domain from an envelope recipient is changed by the resolver as a result of the **qualify\_single** or **search\_parents** options, Exim rewrites the corresponding address in the message's header lines unless **rewrite\_headers** is set false. Exim then re-routes the address, using the full domain.

These two options affect only the DNS lookup that takes place inside the router for the domain of the address that is being routed. They do not affect lookups such as that implied by

```
domains = @mx_any
```

that may happen while processing a router precondition before the router is entered. No widening ever takes place for these lookups.

## 18. The ipliteral router

This router has no private options. Unless it is being used purely for verification (see **verify\_only**) a transport is required to be defined by the generic **transport** option. The router accepts the address if its domain part takes the form of an RFC 2822 domain literal. For example, the *ipliteral* router handles the address

```
root@[192.168.1.1]
```

by setting up delivery to the host with that IP address. IPv4 domain literals consist of an IPv4 address enclosed in square brackets. IPv6 domain literals are similar, but the address is preceded by `ipv6:`. For example:

```
postmaster@[ipv6:fe80::a00:20ff:fe86:a061.5678]
```

Exim allows `ipv4:` before IPv4 addresses, for consistency, and on the grounds that sooner or later somebody will try it.

If the IP address matches something in **ignore\_target\_hosts**, the router declines. If an IP literal turns out to refer to the local host, the generic **self** option determines what happens.

The RFCs require support for domain literals; however, their use is controversial in today's Internet. If you want to use this router, you must also set the main configuration option **allow\_domain\_literals**. Otherwise, Exim will not recognize the domain literal syntax in addresses.



## 19. The iplookup router

The *iplookup* router was written to fulfil a specific requirement in Cambridge University (which in fact no longer exists). For this reason, it is not included in the binary of Exim by default. If you want to include it, you must set

```
ROUTER_IPLOOKUP=yes
```

in your *Local/Makefile* configuration file.

The *iplookup* router routes an address by sending it over a TCP or UDP connection to one or more specific hosts. The host can then return the same or a different address – in effect rewriting the recipient address in the message’s envelope. The new address is then passed on to subsequent routers. If this process fails, the address can be passed on to other routers, or delivery can be deferred. Since *iplookup* is just a rewriting router, a transport must not be specified for it.

<b>hosts</b>	Use: <i>iplookup</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------	----------------------	---------------------	-----------------------

This option must be supplied. Its value is a colon-separated list of host names. The hosts are looked up using *gethostbyname()* (or *getipnodebyname()* when available) and are tried in order until one responds to the query. If none respond, what happens is controlled by **optional**.

<b>optional</b>	Use: <i>iplookup</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------	----------------------	----------------------	-----------------------

If **optional** is true, if no response is obtained from any host, the address is passed to the next router, overriding **no\_more**. If **optional** is false, delivery to the address is deferred.

<b>port</b>	Use: <i>iplookup</i>	Type: <i>integer</i>	Default: <i>0</i>
-------------	----------------------	----------------------	-------------------

This option must be supplied. It specifies the port number for the TCP or UDP call.

<b>protocol</b>	Use: <i>iplookup</i>	Type: <i>string</i>	Default: <i>udp</i>
-----------------	----------------------	---------------------	---------------------

This option can be set to “udp” or “tcp” to specify which of the two protocols is to be used.

<b>query</b>	Use: <i>iplookup</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
--------------	----------------------	----------------------------------	---------------------------

This defines the content of the query that is sent to the remote hosts. The default value is:

```
$local_part@$domain $local_part@$domain
```

The repetition serves as a way of checking that a response is to the correct query in the default case (see **response\_pattern** below).

<b>reroute</b>	Use: <i>iplookup</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
----------------	----------------------	----------------------------------	-----------------------

If this option is not set, the rerouted address is precisely the byte string returned by the remote host, up to the first white space, if any. If set, the string is expanded to form the rerouted address. It can include parts matched in the response by **response\_pattern** by means of numeric variables such as *\$1*, *\$2*, etc. The variable *\$0* refers to the entire input string, whether or not a pattern is in use. In all cases, the rerouted address must end up in the form *local\_part@domain*.

<b>response_pattern</b>	Use: <i>iplookup</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------------	----------------------	---------------------	-----------------------

This option can be set to a regular expression that is applied to the string returned from the remote host. If the pattern does not match the response, the router declines. If **response\_pattern** is not set, no checking of the response is done, unless the query was defaulted, in which case there is a check that the text returned after the first white space is the original address. This checks that the answer that has been received is in response to the correct question. For example, if the response is just a new domain, the following could be used:

```
response_pattern = ^([^\@]+)$
reroute = $local_part@$1
```

<b>timeout</b>	Use: <i>iplookup</i>	Type: <i>time</i>	Default: <i>5s</i>
----------------	----------------------	-------------------	--------------------

This specifies the amount of time to wait for a response from the remote machine. The same timeout is used for the *connect()* function for a TCP call. It does not apply to UDP.

## 20. The manualroute router

The *manualroute* router is so-called because it provides a way of manually routing an address according to its domain. It is mainly used when you want to route addresses to remote hosts according to your own rules, bypassing the normal DNS routing that looks up MX records. However, *manualroute* can also route to local transports, a facility that may be useful if you want to save messages for dial-in hosts in local files.

The *manualroute* router compares a list of domain patterns with the domain it is trying to route. If there is no match, the router declines. Each pattern has associated with it a list of hosts and some other optional data, which may include a transport. The combination of a pattern and its data is called a “routing rule”. For patterns that do not have an associated transport, the generic **transport** option must specify a transport, unless the router is being used purely for verification (see **verify\_only**).

In the case of verification, matching the domain pattern is sufficient for the router to accept the address. When actually routing an address for delivery, an address that matches a domain pattern is queued for the associated transport. If the transport is not a local one, a host list must be associated with the pattern; IP addresses are looked up for the hosts, and these are passed to the transport along with the mail address. For local transports, a host list is optional. If it is present, it is passed in *\$host* as a single text string.

The list of routing rules can be provided as an inline string in **route\_list**, or the data can be obtained by looking up the domain in a file or database by setting **route\_data**. Only one of these settings may appear in any one instance of *manualroute*. The format of routing rules is described below, following the list of private options.

### 20.1 Private options for manualroute

The private options for the *manualroute* router are as follows:

<b>host_all_ignored</b>	Use: <i>manualroute</i>	Type: <i>string</i>	Default: <i>defer</i>
-------------------------	-------------------------	---------------------	-----------------------

See **host\_find\_failed**.

<b>host_find_failed</b>	Use: <i>manualroute</i>	Type: <i>string</i>	Default: <i>freeze</i>
-------------------------	-------------------------	---------------------	------------------------

This option controls what happens when *manualroute* tries to find an IP address for a host, and the host does not exist. The option can be set to one of the following values:

```
decline
defer
fail
freeze
ignore
pass
```

The default (“freeze”) assumes that this state is a serious configuration error. The difference between “pass” and “decline” is that the former forces the address to be passed to the next router (or the router defined by **pass\_router**), overriding **no\_more**, whereas the latter passes the address to the next router only if **more** is true.

The value “ignore” causes Exim to completely ignore a host whose IP address cannot be found. If all the hosts in the list are ignored, the behaviour is controlled by the **host\_all\_ignored** option. This takes the same values as **host\_find\_failed**, except that it cannot be set to “ignore”.

The **host\_find\_failed** option applies only to a definite “does not exist” state; if a host lookup gets a temporary error, delivery is deferred unless the generic **pass\_on\_timeout** option is set.

<b>hosts_randomize</b>	Use: <i>manualroute</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	-------------------------	----------------------	-----------------------

If this option is set, the order of the items in a host list in a routing rule is randomized each time the list is used, unless an option in the routing rule overrides (see below). Randomizing the order of a host list can be used to do crude load sharing. However, if more than one mail address is routed by the same router to the same host list, the host lists are considered to be the same (even though they may be randomized into different orders) for the purpose of deciding whether to batch the deliveries into a single SMTP transaction.

When **hosts\_randomize** is true, a host list may be split into groups whose order is separately randomized. This makes it possible to set up MX-like behaviour. The boundaries between groups are indicated by an item that is just + in the host list. For example:

```
route_list = * host1:host2:host3:+:host4:host5
```

The order of the first three hosts and the order of the last two hosts is randomized for each use, but the first three always end up before the last two. If **hosts\_randomize** is not set, a + item in the list is ignored. If a randomized host list is passed to an *smtp* transport that also has **hosts\_randomize set**, the list is not re-randomized.

<b>route_data</b>	Use: <i>manualroute</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------	-------------------------	----------------------	-----------------------

If this option is set, it must expand to yield the data part of a routing rule. Typically, the expansion string includes a lookup based on the domain. For example:

```
route_data = ${lookup{$domain}dbm{/etc/routes}}
```

If the expansion is forced to fail, or the result is an empty string, the router declines. Other kinds of expansion failure cause delivery to be deferred.

<b>route_list</b>	Use: <i>manualroute</i>	Type: <i>string list</i>	Default: <i>unset</i>
-------------------	-------------------------	--------------------------	-----------------------

This string is a list of routing rules, in the form defined below. Note that, unlike most string lists, the items are separated by semicolons. This is so that they may contain colon-separated host lists.

<b>same_domain_copy_routing</b>	Use: <i>manualroute</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------------	-------------------------	----------------------	-----------------------

Addresses with the same domain are normally routed by the *manualroute* router to the same list of hosts. However, this cannot be presumed, because the router options and preconditions may refer to the local part of the address. By default, therefore, Exim routes each address in a message independently. DNS servers run caches, so repeated DNS lookups are not normally expensive, and in any case, personal messages rarely have more than a few recipients.

If you are running mailing lists with large numbers of subscribers at the same domain, and you are using a *manualroute* router which is independent of the local part, you can set **same\_domain\_copy\_routing** to bypass repeated DNS lookups for identical domains in one message. In this case, when *manualroute* routes an address to a remote transport, any other unrouted addresses in the message that have the same domain are automatically given the same routing without processing them independently. However, this is only done if **headers\_add** and **headers\_remove** are unset.

## 20.2 Routing rules in route\_list

The value of **route\_list** is a string consisting of a sequence of routing rules, separated by semicolons. If a semicolon is needed in a rule, it can be entered as two semicolons. Alternatively, the list separator can be changed as described (for colon-separated lists) in section 6.20. Empty rules are ignored. The format of each rule is

```
<domain pattern> <list of hosts> <options>
```

The following example contains two rules, each with a simple domain pattern and no options:

```
route_list = \
    dict.ref.example mail-1.ref.example:mail-2.ref.example ; \
    thes.ref.example mail-3.ref.example:mail-4.ref.example
```

The three parts of a rule are separated by white space. The pattern and the list of hosts can be enclosed in quotes if necessary, and if they are, the usual quoting rules apply. Each rule in a **route\_list** must start with a single domain pattern, which is the only mandatory item in the rule. The pattern is in the same format as one item in a domain list (see section 10.8), except that it may not be the name of an interpolated file. That is, it may be wildcarded, or a regular expression, or a file or database lookup (with semicolons doubled, because of the use of semicolon as a separator in a **route\_list**).

The rules in **route\_list** are searched in order until one of the patterns matches the domain that is being routed. The list of hosts and then options are then used as described below. If there is no match, the router declines. When **route\_list** is set, **route\_data** must not be set.

## 20.3 Routing rules in route\_data

The use of **route\_list** is convenient when there are only a small number of routing rules. For larger numbers, it is easier to use a file or database to hold the routing information, and use the **route\_data** option instead. The value of **route\_data** is a list of hosts, followed by (optional) options. Most commonly, **route\_data** is set as a string that contains an expansion lookup. For example, suppose we place two routing rules in a file like this:

```
dict.ref.example: mail-1.ref.example:mail-2.ref.example
thes.ref.example: mail-3.ref.example:mail-4.ref.example
```

This data can be accessed by setting

```
route_data = ${lookup{$domain}lsearch{/the/file/name}}
```

Failure of the lookup results in an empty string, causing the router to decline. However, you do not have to use a lookup in **route\_data**. The only requirement is that the result of expanding the string is a list of hosts, possibly followed by options, separated by white space. The list of hosts must be enclosed in quotes if it contains white space.

## 20.4 Format of the list of hosts

A list of hosts, whether obtained via **route\_data** or **route\_list**, is always separately expanded before use. If the expansion fails, the router declines. The result of the expansion must be a colon-separated list of names and/or IP addresses, optionally also including ports. The format of each item in the list is described in the next section. The list separator can be changed as described in section 6.20.

If the list of hosts was obtained from a **route\_list** item, the following variables are set during its expansion:

- If the domain was matched against a regular expression, the numeric variables *\$1*, *\$2*, etc. may be set. For example:

```
route_list = ^domain(\d+) host-$1.text.example
```

- *\$0* is always set to the entire domain.
- *\$1* is also set when partial matching is done in a file lookup.
- If the pattern that matched the domain was a lookup item, the data that was looked up is available in the expansion variable *\$value*. For example:

```
route_list = lsearch;;/some/file.routes $value
```

Note the doubling of the semicolon in the pattern that is necessary because semicolon is the default route list separator.

## 20.5 Format of one host item

Each item in the list of hosts is either a host name or an IP address, optionally with an attached port number. When no port is given, an IP address is not enclosed in brackets. When a port is specified, it overrides the port specification on the transport. The port is separated from the name or address by a colon. This leads to some complications:

- Because colon is the default separator for the list of hosts, either the colon that specifies a port must be doubled, or the list separator must be changed. The following two examples have the same effect:

```
route_list = * "host1.tld::1225 : host2.tld::1226"
route_list = * "<+ host1.tld:1225 + host2.tld:1226"
```

- When IPv6 addresses are involved, it gets worse, because they contain colons of their own. To make this case easier, it is permitted to enclose an IP address (either v4 or v6) in square brackets if a port number follows. For example:

```
route_list = * "</ [10.1.1.1]:1225 / [::1]:1226"
```

## 20.6 How the list of hosts is used

When an address is routed to an *smtp* transport by *manualroute*, each of the hosts is tried, in the order specified, when carrying out the SMTP delivery. However, the order can be changed by setting the **hosts\_randomize** option, either on the router (see section 20.1 above), or on the transport.

Hosts may be listed by name or by IP address. An unadorned name in the list of hosts is interpreted as a host name. A name that is followed by **/MX** is interpreted as an indirection to a sublist of hosts obtained by looking up MX records in the DNS. For example:

```
route_list = * x.y.z:p.q.r/MX:e.f.g
```

If this feature is used with a port specifier, the port must come last. For example:

```
route_list = * dom1.tld/mx::1225
```

If the **hosts\_randomize** option is set, the order of the items in the list is randomized before any lookups are done. Exim then scans the list; for any name that is not followed by **/MX** it looks up an IP address. If this turns out to be an interface on the local host and the item is not the first in the list, Exim discards it and any subsequent items. If it is the first item, what happens is controlled by the **self** option of the router.

A name on the list that is followed by **/MX** is replaced with the list of hosts obtained by looking up MX records for the name. This is always a DNS lookup; the **bydns** and **byname** options (see section 20.7 below) are not relevant here. The order of these hosts is determined by the preference values in the MX records, according to the usual rules. Because randomizing happens before the MX lookup, it does not affect the order that is defined by MX preferences.

If the local host is present in the sublist obtained from MX records, but is not the most preferred host in that list, it and any equally or less preferred hosts are removed before the sublist is inserted into the main list.

If the local host is the most preferred host in the MX list, what happens depends on where in the original list of hosts the **/MX** item appears. If it is not the first item (that is, there are previous hosts in the main list), Exim discards this name and any subsequent items in the main list.

If the MX item is first in the list of hosts, and the local host is the most preferred host, what happens is controlled by the **self** option of the router.

DNS failures when lookup up the MX records are treated in the same way as DNS failures when looking up IP addresses: **pass\_on\_timeout** and **host\_find\_failed** are used when relevant.

The generic **ignore\_target\_hosts** option applies to all hosts in the list, whether obtained from an MX lookup or not.

## 20.7 How the options are used

The options are a sequence of words; in practice no more than three are ever present. One of the words can be the name of a transport; this overrides the **transport** option on the router for this particular routing rule only. The other words (if present) control randomization of the list of hosts on a per-rule basis, and how the IP addresses of the hosts are to be found when routing to a remote transport. These options are as follows:

- **randomize**: randomize the order of the hosts in this list, overriding the setting of **hosts\_randomize** for this routing rule only.
- **no\_randomize**: do not randomize the order of the hosts in this list, overriding the setting of **hosts\_randomize** for this routing rule only.
- **byname**: use *getipnodebyname()* (*gethostbyname()* on older systems) to find IP addresses. This function may ultimately cause a DNS lookup, but it may also look in */etc/hosts* or other sources of information.
- **bydns**: look up address records for the hosts directly in the DNS; fail if no address records are found. If there is a temporary DNS error (such as a timeout), delivery is deferred.

For example:

```
route_list = domain1  host1:host2:host3  randomize bydns;\
              domain2  host4:host5
```

If neither **byname** nor **bydns** is given, Exim behaves as follows: First, a DNS lookup is done. If this yields anything other than **HOST\_NOT\_FOUND**, that result is used. Otherwise, Exim goes on to try a call to *getipnodebyname()* or *gethostbyname()*, and the result of the lookup is the result of that call.

**Warning:** It has been discovered that on some systems, if a DNS lookup called via *getipnodebyname()* times out, **HOST\_NOT\_FOUND** is returned instead of **TRY\_AGAIN**. That is why the default action is to try a DNS lookup first. Only if that gives a definite “no such host” is the local function called.

**Compatibility:** From Exim 4.85 until fixed for 4.90, there was an inadvertent constraint that a transport name as an option had to be the last option specified.

If no IP address for a host can be found, what happens is controlled by the **host\_find\_failed** option.

When an address is routed to a local transport, IP addresses are not looked up. The host list is passed to the transport in the *\$host* variable.

## 20.8 Manualroute examples

In some of the examples that follow, the presence of the **remote\_smtp** transport, as defined in the default configuration file, is assumed:

- The *manualroute* router can be used to forward all external mail to a *smart host*. If you have set up, in the main part of the configuration, a named domain list that contains your local domains, for example:

```
domainlist local_domains = my.domain.example
```

You can arrange for all other domains to be routed to a smart host by making your first router something like this:

```
smart_route:
  driver = manualroute
  domains = !+local_domains
  transport = remote_smtp
  route_list = * smarthost.ref.example
```

This causes all non-local addresses to be sent to the single host *smarthost.ref.example*. If a colon-separated list of smart hosts is given, they are tried in order (but you can use **hosts\_randomize** to vary the order each time). Another way of configuring the same thing is this:

```
smart_route:
    driver = manualroute
    transport = remote_smtp
    route_list = !+local_domains    smarthost.ref.example
```

There is no difference in behaviour between these two routers as they stand. However, they behave differently if **no\_more** is added to them. In the first example, the router is skipped if the domain does not match the **domains** precondition; the following router is always tried. If the router runs, it always matches the domain and so can never decline. Therefore, **no\_more** would have no effect. In the second case, the router is never skipped; it always runs. However, if it doesn't match the domain, it declines. In this case **no\_more** would prevent subsequent routers from running.

- A *mail hub* is a host which receives mail for a number of domains via MX records in the DNS and delivers it via its own private routing mechanism. Often the final destinations are behind a firewall, with the mail hub being the one machine that can connect to machines both inside and outside the firewall. The *manualroute* router is usually used on a mail hub to route incoming messages to the correct hosts. For a small number of domains, the routing can be inline, using the **route\_list** option, but for a larger number a file or database lookup is easier to manage.

If the domain names are in fact the names of the machines to which the mail is to be sent by the mail hub, the configuration can be quite simple. For example:

```
hub_route:
    driver = manualroute
    transport = remote_smtp
    route_list = *.rhodes.tvs.example    $domain
```

This configuration routes domains that match *\*.rhodes.tvs.example* to hosts whose names are the same as the mail domains. A similar approach can be taken if the host name can be obtained from the domain name by a string manipulation that the expansion facilities can handle. Otherwise, a lookup based on the domain can be used to find the host:

```
through_firewall:
    driver = manualroute
    transport = remote_smtp
    route_data = ${lookup {$domain} cdb {/internal/host/routes}}
```

The result of the lookup must be the name or IP address of the host (or hosts) to which the address is to be routed. If the lookup fails, the route data is empty, causing the router to decline. The address then passes to the next router.

- You can use *manualroute* to deliver messages to pipes or files in batched SMTP format for onward transportation by some other means. This is one way of storing mail for a dial-up host when it is not connected. The route list entry can be as simple as a single domain name in a configuration like this:

```
save_in_file:
    driver = manualroute
    transport = batchsmtp_appendfile
    route_list = saved.domain.example
```

though often a pattern is used to pick up more than one domain. If there are several domains or groups of domains with different transport requirements, different transports can be listed in the routing information:

```
save_in_file:
    driver = manualroute
    route_list = \
        *.saved.domain1.example    $domain    batch_appendfile; \
        *.saved.domain2.example    \
        ${lookup{$domain}dbm{/domain2/hosts}{$value}fail} \
        batch_pipe
```



The first of these just passes the domain in the *\$host* variable, which doesn't achieve much (since it is also in *\$domain*), but the second does a file lookup to find a value to pass, causing the router to decline to handle the address if the lookup fails.

- Routing mail directly to UUCP software is a specific case of the use of *manualroute* in a gateway to another mail environment. This is an example of one way it can be done:

```
# Transport
uucp:
  driver = pipe
  user = nobody
  command = /usr/local/bin/uux -r - \
    ${substr_-5:$host}!rmail ${local_part}
  return_fail_output = true

# Router
uucphost:
  transport = uucp
  driver = manualroute
  route_data = \
    ${lookup{$domain}lsearch{/usr/local/exim/uucphosts}}
```

The file */usr/local/exim/uucphosts* contains entries like

```
darksite.ethereal.example:          darksite.UUCP
```

It can be set up more simply without adding and removing “.UUCP” but this way makes clear the distinction between the domain name *darksite.ethereal.example* and the UUCP host name *darksite*.

## 21. The queryprogram router

The *queryprogram* router routes an address by running an external command and acting on its output. This is an expensive way to route, and is intended mainly for use in lightly-loaded systems, or for performing experiments. However, if it is possible to use the precondition options (**domains**, **local\_parts**, etc) to skip this router for most addresses, it could sensibly be used in special cases, even on a busy host. There are the following private options:

<b>command</b>	Use: <i>queryprogram</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
----------------	--------------------------	----------------------------------	-----------------------

This option must be set. It specifies the command that is to be run. The command is split up into a command name and arguments, and then each is expanded separately (exactly as for a *pipe* transport, described in chapter 29).

<b>command_group</b>	Use: <i>queryprogram</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------	--------------------------	---------------------	-----------------------

This option specifies a gid to be set when running the command while routing an address for deliver. It must be set if **command\_user** specifies a numerical uid. If it begins with a digit, it is interpreted as the numerical value of the gid. Otherwise it is looked up using *getgrnam()*.

<b>command_user</b>	Use: <i>queryprogram</i>	Type: <i>string</i>	Default: <i>unset</i>
---------------------	--------------------------	---------------------	-----------------------

This option must be set. It specifies the uid which is set when running the command while routing an address for delivery. If the value begins with a digit, it is interpreted as the numerical value of the uid. Otherwise, it is looked up using *getpwnam()* to obtain a value for the uid and, if **command\_group** is not set, a value for the gid also.

**Warning:** Changing uid and gid is possible only when Exim is running as root, which it does during a normal delivery in a conventional configuration. However, when an address is being verified during message reception, Exim is usually running as the Exim user, not as root. If the *queryprogram* router is called from a non-root process, Exim cannot change uid or gid before running the command. In this circumstance the command runs under the current uid and gid.

<b>current_directory</b>	Use: <i>queryprogram</i>	Type: <i>string</i>	Default: <i>/</i>
--------------------------	--------------------------	---------------------	-------------------

This option specifies an absolute path which is made the current directory before running the command.

<b>timeout</b>	Use: <i>queryprogram</i>	Type: <i>time</i>	Default: <i>1h</i>
----------------	--------------------------	-------------------	--------------------

If the command does not complete within the timeout period, its process group is killed and the message is frozen. A value of zero time specifies no timeout.

The standard output of the command is connected to a pipe, which is read when the command terminates. It should consist of a single line of output, containing up to five fields, separated by white space. The maximum length of the line is 1023 characters. Longer lines are silently truncated. The first field is one of the following words (case-insensitive):

- *Accept*: routing succeeded; the remaining fields specify what to do (see below).
- *Decline*: the router declines; pass the address to the next router, unless **no\_more** is set.
- *Fail*: routing failed; do not pass the address to any more routers. Any subsequent text on the line is an error message. If the router is run as part of address verification during an incoming SMTP message, the message is included in the SMTP response.

- *Defer*: routing could not be completed at this time; try again later. Any subsequent text on the line is an error message which is logged. It is not included in any SMTP response.
- *Freeze*: the same as *defer*, except that the message is frozen.
- *Pass*: pass the address to the next router (or the router specified by **pass\_router**), overriding **no\_more**.
- *Redirect*: the message is redirected. The remainder of the line is a list of new addresses, which are routed independently, starting with the first router, or the router specified by **redirect\_router**, if set.

When the first word is *accept*, the remainder of the line consists of a number of keyed data values, as follows (split into two lines here, to fit on the page):

```
ACCEPT TRANSPORT=<transport> HOSTS=<list of hosts>
LOOKUP=byname|bydns DATA=<text>
```

The data items can be given in any order, and all are optional. If no transport is included, the transport specified by the generic **transport** option is used. The list of hosts and the lookup type are needed only if the transport is an *smtp* transport that does not itself supply a list of hosts.

The format of the list of hosts is the same as for the *manualroute* router. As well as host names and IP addresses with optional port numbers, as described in section 20.5, it may contain names followed by /MX to specify sublists of hosts that are obtained by looking up MX records (see section 20.6).

If the lookup type is not specified, Exim behaves as follows when trying to find an IP address for each host: First, a DNS lookup is done. If this yields anything other than HOST\_NOT\_FOUND, that result is used. Otherwise, Exim goes on to try a call to *getipnodebyname()* or *gethostbyname()*, and the result of the lookup is the result of that call.

If the DATA field is set, its value is placed in the *\$address\_data* variable. For example, this return line

```
accept hosts=x1.y.example:x2.y.example data="rule1"
```

routes the address to the default transport, passing a list of two hosts. When the transport runs, the string "rule1" is in *\$address\_data*.

## 22. The redirect router

The *redirect* router handles several kinds of address redirection. Its most common uses are for resolving local part aliases from a central alias file (usually called */etc/aliases*) and for handling users' personal *.forward* files, but it has many other potential uses. The incoming address can be redirected in several different ways:

- It can be replaced by one or more new addresses which are themselves routed independently.
- It can be routed to be delivered to a given file or directory.
- It can be routed to be delivered to a specified pipe command.
- It can cause an automatic reply to be generated.
- It can be forced to fail, optionally with a custom error message.
- It can be temporarily deferred, optionally with a custom message.
- It can be discarded.

The generic **transport** option must not be set for *redirect* routers. However, there are some private options which define transports for delivery to files and pipes, and for generating autoreplies. See the **file\_transport**, **pipe\_transport** and **reply\_transport** descriptions below.

If success DSNs have been requested redirection triggers one and the DSN options are not passed any further.

### 22.1 Redirection data

The router operates by interpreting a text string which it obtains either by expanding the contents of the **data** option, or by reading the entire contents of a file whose name is given in the **file** option. These two options are mutually exclusive. The first is commonly used for handling system aliases, in a configuration like this:

```
system_aliases:
    driver = redirect
    data = ${lookup{$local_part}lsearch{/etc/aliases}}
```

If the lookup fails, the expanded string in this example is empty. When the expansion of **data** results in an empty string, the router declines. A forced expansion failure also causes the router to decline; other expansion failures cause delivery to be deferred.

A configuration using **file** is commonly used for handling users' *.forward* files, like this:

```
userforward:
    driver = redirect
    check_local_user
    file = $home/.forward
    no_verify
```

If the file does not exist, or causes no action to be taken (for example, it is empty or consists only of comments), the router declines. **Warning:** This is not the case when the file contains syntactically valid items that happen to yield empty addresses, for example, items containing only RFC 2822 address comments.

### 22.2 Forward files and address verification

It is usual to set **no\_verify** on *redirect* routers which handle users' *.forward* files, as in the example above. There are two reasons for this:

- When Exim is receiving an incoming SMTP message from a remote host, it is running under the Exim uid, not as root. Exim is unable to change uid to read the file as the user, and it may not be able to read it as the Exim user. So in practice the router may not be able to operate.

- However, even when the router can operate, the existence of a *.forward* file is unimportant when verifying an address. What should be checked is whether the local part is a valid user name or not. Cutting out the redirection processing saves some resources.

## 22.3 Interpreting redirection data

The contents of the data string, whether obtained from **data** or **file**, can be interpreted in two different ways:

- If the **allow\_filter** option is set true, and the data begins with the text “#Exim filter” or “#Sieve filter”, it is interpreted as a list of *filtering* instructions in the form of an Exim or Sieve filter file, respectively. Details of the syntax and semantics of filter files are described in a separate document entitled *Exim’s interfaces to mail filtering*; this document is intended for use by end users.
- Otherwise, the data must be a comma-separated list of redirection items, as described in the next section.

When a message is redirected to a file (a “mail folder”), the file name given in a non-filter redirection list must always be an absolute path. A filter may generate a relative path – how this is handled depends on the transport’s configuration. See section 26.1 for a discussion of this issue for the *appendfile* transport.

## 22.4 Items in a non-filter redirection list

When the redirection data is not an Exim or Sieve filter, for example, if it comes from a conventional alias or forward file, it consists of a list of addresses, file names, pipe commands, or certain special items (see section 22.6 below). The special items can be individually enabled or disabled by means of options whose names begin with **allow\_** or **forbid\_**, depending on their default values. The items in the list are separated by commas or newlines. If a comma is required in an item, the entire item must be enclosed in double quotes.

Lines starting with a # character are comments, and are ignored, and # may also appear following a comma, in which case everything between the # and the next newline character is ignored.

If an item is entirely enclosed in double quotes, these are removed. Otherwise double quotes are retained because some forms of mail address require their use (but never to enclose the entire address). In the following description, “item” refers to what remains after any surrounding double quotes have been removed.

**Warning:** If you use an Exim expansion to construct a redirection address, and the expansion contains a reference to *\$local\_part*, you should make use of the **quote\_local\_part** expansion operator, in case the local part contains special characters. For example, to redirect all mail for the domain *obsolete.example*, retaining the existing local part, you could use this setting:

```
data = ${quote_local_part:$local_part}@newdomain.example
```

## 22.5 Redirecting to a local mailbox

A redirection item may safely be the same as the address currently under consideration. This does not cause a routing loop, because a router is automatically skipped if any ancestor of the address that is being processed is the same as the current address and was processed by the current router. Such an address is therefore passed to the following routers, so it is handled as if there were no redirection. When making this loop-avoidance test, the complete local part, including any prefix or suffix, is used.

Specifying the same local part without a domain is a common usage in personal filter files when the user wants to have messages delivered to the local mailbox and also forwarded elsewhere. For example, the user whose login is *cleo* might have a *.forward* file containing this:

```
cleo, cleopatra@egypt.example
```

For compatibility with other MTAs, such unqualified local parts may be preceded by “\”, but this is not a requirement for loop prevention. However, it does make a difference if more than one domain is being handled synonymously.

If an item begins with “\” and the rest of the item parses as a valid RFC 2822 address that does not include a domain, the item is qualified using the domain of the incoming address. In the absence of a leading “\”, unqualified addresses are qualified using the value in **qualify\_recipient**, but you can force the incoming domain to be used by setting **qualify\_preserve\_domain**.

Care must be taken if there are alias names for local users. Consider an MTA handling a single local domain where the system alias file contains:

```
Sam.Reman: spqr
```

Now suppose that Sam (whose login id is *spqr*) wants to save copies of messages in the local mailbox, and also forward copies elsewhere. He creates this forward file:

```
Sam.Reman, spqr@reme.elsewhere.example
```

With these settings, an incoming message addressed to *Sam.Reman* fails. The *redirect* router for system aliases does not process *Sam.Reman* the second time round, because it has previously routed it, and the following routers presumably cannot handle the alias. The forward file should really contain

```
spqr, spqr@reme.elsewhere.example
```

but because this is such a common error, the **check\_ancestor** option (see below) exists to provide a way to get round it. This is normally set on a *redirect* router that is handling users’ *.forward* files.

## 22.6 Special items in redirection lists

In addition to addresses, the following types of item may appear in redirection lists (that is, in non-filter redirection data):

- An item is treated as a pipe command if it begins with “|” and does not parse as a valid RFC 2822 address that includes a domain. A transport for running the command must be specified by the **pipe\_transport** option. Normally, either the router or the transport specifies a user and a group under which to run the delivery. The default is to use the Exim user and group.

Single or double quotes can be used for enclosing the individual arguments of the pipe command; no interpretation of escapes is done for single quotes. If the command contains a comma character, it is necessary to put the whole item in double quotes, for example:

```
"|/some/command ready,steady,go"
```

since items in redirection lists are terminated by commas. Do not, however, quote just the command. An item such as

```
|"/some/command ready,steady,go"
```

is interpreted as a pipe with a rather strange command name, and no arguments.

Note that the above example assumes that the text comes from a lookup source of some sort, so that the quotes are part of the data. If composing a redirect router with a **data** option directly specifying this command, the quotes will be used by the configuration parser to define the extent of one string, but will not be passed down into the redirect router itself. There are two main approaches to get around this: escape quotes to be part of the data itself, or avoid using this mechanism and instead create a custom transport with the **command** option set and reference that transport from an **accept** router.

- An item is interpreted as a path name if it begins with “/” and does not parse as a valid RFC 2822 address that includes a domain. For example,

```
/home/world/minbari
```

is treated as a file name, but

```
/s=molari/o=babylon/@x400gate.way
```

is treated as an address. For a file name, a transport must be specified using the **file\_transport** option. However, if the generated path name ends with a forward slash character, it is interpreted as a directory name rather than a file name, and **directory\_transport** is used instead.

Normally, either the router or the transport specifies a user and a group under which to run the delivery. The default is to use the Exim user and group.

However, if a redirection item is the path */dev/null*, delivery to it is bypassed at a high level, and the log entry shows “\*\*bypassed\*\*” instead of a transport name. In this case the user and group are not used.

- If an item is of the form

```
:include:<path name>
```

a list of further items is taken from the given file and included at that point. **Note:** Such a file can not be a filter file; it is just an out-of-line addition to the list. The items in the included list are separated by commas or newlines and are not subject to expansion. If this is the first item in an alias list in an *lsearch* file, a colon must be used to terminate the alias name. This example is incorrect:

```
list1      :include:/opt/lists/list1
```

It must be given as

```
list1:     :include:/opt/lists/list1
```

- Sometimes you want to throw away mail to a particular local part. Making the **data** option expand to an empty string does not work, because that causes the router to decline. Instead, the alias item

```
:blackhole:
```

can be used. It does what its name implies. No delivery is done, and no error message is generated. This has the same effect as specifying */dev/null* as a destination, but it can be independently disabled.

**Warning:** If *:blackhole:* appears anywhere in a redirection list, no delivery is done for the original local part, even if other redirection items are present. If you are generating a multi-item list (for example, by reading a database) and need the ability to provide a no-op item, you must use */dev/null*.

- An attempt to deliver a particular address can be deferred or forced to fail by redirection items of the form

```
:defer:  
:fail:
```

respectively. When a redirection list contains such an item, it applies to the entire redirection; any other items in the list are ignored. Any text following *:fail:* or *:defer:* is placed in the error text associated with the failure. For example, an alias file might contain:

```
X.Employee:  :fail: Gone away, no forwarding address
```

In the case of an address that is being verified from an ACL or as the subject of a VRFY command, the text is included in the SMTP error response by default. The text is not included in the response to an EXPN command. In non-SMTP cases the text is included in the error message that Exim generates.

By default, Exim sends a 451 SMTP code for a *:defer:*, and 550 for *:fail:*. However, if the message starts with three digits followed by a space, optionally followed by an extended code of the form *n.n.n*, also followed by a space, and the very first digit is the same as the default error code, the code from the message is used instead. If the very first digit is incorrect, a panic error is logged, and the default code is used. You can suppress the use of the supplied code in a redirect router by setting the **forbid\_smtp\_code** option true. In this case, any SMTP code is quietly ignored.

In an ACL, an explicitly provided message overrides the default, but the default message is available in the variable `$acl_verify_message` and can therefore be included in a custom message if this is desired.

Normally the error text is the rest of the redirection list – a comma does not terminate it – but a newline does act as a terminator. Newlines are not normally present in alias expansions. In *lsearch* lookups they are removed as part of the continuation process, but they may exist in other kinds of lookup and in *:include:* files.

During routing for message delivery (as opposed to verification), a redirection containing *:fail:* causes an immediate failure of the incoming address, whereas *:defer:* causes the message to remain on the queue so that a subsequent delivery attempt can happen at a later time. If an address is deferred for too long, it will ultimately fail, because the normal retry rules still apply.

- Sometimes it is useful to use a single-key search type with a default (see chapter 9) to look up aliases. However, there may be a need for exceptions to the default. These can be handled by aliasing them to *:unknown:*. This differs from *:fail:* in that it causes the *redirect* router to decline, whereas *:fail:* forces routing to fail. A lookup which results in an empty redirection list has the same effect.

## 22.7 Duplicate addresses

Exim removes duplicate addresses from the list to which it is delivering, so as to deliver just one copy to each address. This does not apply to deliveries routed to pipes by different immediate parent addresses, but an indirect aliasing scheme of the type

```
pipe:          | /some/command $local_part
localpart1:    pipe
localpart2:    pipe
```

does not work with a message that is addressed to both local parts, because when the second is aliased to the intermediate local part “pipe” it gets discarded as being the same as a previously handled address. However, a scheme such as

```
localpart1:    | /some/command $local_part
localpart2:    | /some/command $local_part
```

does result in two different pipe deliveries, because the immediate parents of the pipes are distinct.

## 22.8 Repeated redirection expansion

When a message cannot be delivered to all of its recipients immediately, leading to two or more delivery attempts, redirection expansion is carried out afresh each time for those addresses whose children were not all previously delivered. If redirection is being used as a mailing list, this can lead to new members of the list receiving copies of old messages. The **one\_time** option can be used to avoid this.

## 22.9 Errors in redirection lists

If **skip\_syntax\_errors** is set, a malformed address that causes a parsing error is skipped, and an entry is written to the main log. This may be useful for mailing lists that are automatically managed. Otherwise, if an error is detected while generating the list of new addresses, the original address is deferred. See also **syntax\_errors\_to**.

## 22.10 Private options for the redirect router

The private options for the *redirect* router are as follows:



<b>allow_defer</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	----------------------	----------------------	-----------------------

Setting this option allows the use of *:defer:* in non-filter redirection data, or the **defer** command in an Exim filter file.

<b>allow_fail</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------	----------------------	----------------------	-----------------------

If this option is true, the *:fail:* item can be used in a redirection list, and the **fail** command may be used in an Exim filter file.

<b>allow_filter</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------	----------------------	----------------------	-----------------------

Setting this option allows Exim to interpret redirection data that starts with “#Exim filter” or “#Sieve filter” as a set of filtering instructions. There are some features of Exim filter files that some administrators may wish to lock out; see the **forbid\_filter\_xxx** options below.

It is also possible to lock out Exim filters or Sieve filters while allowing the other type; see **forbid\_exim\_filter** and **forbid\_sieve\_filter**.

The filter is run using the uid and gid set by the generic **user** and **group** options. These take their defaults from the password data if **check\_local\_user** is set, so in the normal case of users’ personal filter files, the filter is run as the relevant user. When **allow\_filter** is set true, Exim insists that either **check\_local\_user** or **user** is set.

<b>allow_freeze</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------	----------------------	----------------------	-----------------------

Setting this option allows the use of the **freeze** command in an Exim filter. This command is more normally encountered in system filters, and is disabled by default for redirection filters because it isn’t something you usually want to let ordinary users do.

<b>check_ancestor</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	----------------------	----------------------	-----------------------

This option is concerned with handling generated addresses that are the same as some address in the list of redirection ancestors of the current address. Although it is turned off by default in the code, it is set in the default configuration file for handling users’ *.forward* files. It is recommended for this use of the *redirect* router.

When **check\_ancestor** is set, if a generated address (including the domain) is the same as any ancestor of the current address, it is replaced by a copy of the current address. This helps in the case where local part A is aliased to B, and B has a *.forward* file pointing back to A. For example, within a single domain, the local part “Joe.Bloggs” is aliased to “jb” and *jb.forward* contains:

```
\Joe.Bloggs, <other item(s)>
```

Without the **check\_ancestor** setting, either local part (“jb” or “joe.bloggs”) gets processed once by each router and so ends up as it was originally. If “jb” is the real mailbox name, mail to “jb” gets delivered (having been turned into “joe.bloggs” by the *.forward* file and back to “jb” by the alias), but mail to “joe.bloggs” fails. Setting **check\_ancestor** on the *redirect* router that handles the *.forward* file prevents it from turning “jb” back into “joe.bloggs” when that was the original address. See also the **repeat\_use** option below.

<b>check_group</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>see below</i>
--------------------	----------------------	----------------------	---------------------------

When the **file** option is used, the group owner of the file is checked only when this option is set. The permitted groups are those listed in the **owngroups** option, together with the user’s default group if **check\_local\_user** is set. If the file has the wrong group, routing is deferred. The default setting for

this option is true if **check\_local\_user** is set and the **modemask** option permits the group write bit, or if the **owngroups** option is set. Otherwise it is false, and no group check occurs.

<b>check_owner</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>see below</i>
--------------------	----------------------	----------------------	---------------------------

When the **file** option is used, the owner of the file is checked only when this option is set. If **check\_local\_user** is set, the local user is permitted; otherwise the owner must be one of those listed in the **owners** option. The default value for this option is true if **check\_local\_user** or **owners** is set. Otherwise the default is false, and no owner check occurs.

<b>data</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------	----------------------	----------------------	-----------------------

This option is mutually exclusive with **file**. One or other of them must be set, but not both. The contents of **data** are expanded, and then used as the list of forwarding items, or as a set of filtering instructions. If the expansion is forced to fail, or the result is an empty string or a string that has no effect (consists entirely of comments), the router declines.

When filtering instructions are used, the string must begin with “#Exim filter”, and all comments in the string, including this initial one, must be terminated with newline characters. For example:

```
data = #Exim filter\n\
      if $h_to: contains Exim then save $home/mail/exim endif
```

If you are reading the data from a database where newlines cannot be included, you can use the *\${sg}* expansion item to turn the escape string of your choice into a newline.

<b>directory_transport</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------------	----------------------	----------------------	-----------------------

A *redirect* router sets up a direct delivery to a directory when a path name ending with a slash is specified as a new “address”. The transport used is specified by this option, which, after expansion, must be the name of a configured transport. This should normally be an *appendfile* transport.

<b>file</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------	----------------------	----------------------	-----------------------

This option specifies the name of a file that contains the redirection data. It is mutually exclusive with the **data** option. The string is expanded before use; if the expansion is forced to fail, the router declines. Other expansion failures cause delivery to be deferred. The result of a successful expansion must be an absolute path. The entire file is read and used as the redirection data. If the data is an empty string or a string that has no effect (consists entirely of comments), the router declines.

If the attempt to open the file fails with a “does not exist” error, Exim runs a check on the containing directory, unless **ignore\_enotdir** is true (see below). If the directory does not appear to exist, delivery is deferred. This can happen when users’ *forward* files are in NFS-mounted directories, and there is a mount problem. If the containing directory does exist, but the file does not, the router declines.

<b>file_transport</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
-----------------------	----------------------	----------------------	-----------------------

A *redirect* router sets up a direct delivery to a file when a path name not ending in a slash is specified as a new “address”. The transport used is specified by this option, which, after expansion, must be the name of a configured transport. This should normally be an *appendfile* transport. When it is running, the file name is in *\$address\_file*.

<b>filter_prepend_home</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>true</i>
----------------------------	----------------------	----------------------	----------------------

When this option is true, if a *save* command in an Exim filter specifies a relative path, and *\$home* is defined, it is automatically prepended to the relative path. If this option is set false, this action does not happen. The relative path is then passed to the transport unmodified.

<b>forbid_blackhole</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------	----------------------	----------------------	-----------------------

If this option is true, the *:blackhole:* item may not appear in a redirection list.

<b>forbid_exim_filter</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	----------------------	----------------------	-----------------------

If this option is set true, only Sieve filters are permitted when **allow\_filter** is true.

<b>forbid_file</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	----------------------	----------------------	-----------------------

If this option is true, this router may not generate a new address that specifies delivery to a local file or directory, either from a filter or from a conventional forward file. This option is forced to be true if **one\_time** is set. It applies to Sieve filters as well as to Exim filters, but if true, it locks out the Sieve's "keep" facility.

<b>forbid_filter_dlfunc</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------------	----------------------	----------------------	-----------------------

If this option is true, string expansions in Exim filters are not allowed to make use of the **dlfunc** expansion facility to run dynamically loaded functions.

<b>forbid_filter_existstest</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------------	----------------------	----------------------	-----------------------

If this option is true, string expansions in Exim filters are not allowed to make use of the **exists** condition or the **stat** expansion item.

<b>forbid_filter_logwrite</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------------	----------------------	----------------------	-----------------------

If this option is true, use of the logging facility in Exim filters is not permitted. Logging is in any case available only if the filter is being run under some unprivileged uid (which is normally the case for ordinary users' *.forward* files).

<b>forbid_filter_lookup</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------------	----------------------	----------------------	-----------------------

If this option is true, string expansions in Exim filter files are not allowed to make use of **lookup** items.

<b>forbid_filter_perl</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	----------------------	----------------------	-----------------------

This option has an effect only if Exim is built with embedded Perl support. If it is true, string expansions in Exim filter files are not allowed to make use of the embedded Perl support.

<b>forbid_filter_readfile</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------------	----------------------	----------------------	-----------------------

If this option is true, string expansions in Exim filter files are not allowed to make use of **readfile** items.

<b>forbid_filter_readsocket</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------------	----------------------	----------------------	-----------------------

If this option is true, string expansions in Exim filter files are not allowed to make use of **readsocket** items.

<b>forbid_filter_reply</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------------	----------------------	----------------------	-----------------------

If this option is true, this router may not generate an automatic reply message. Automatic replies can be generated only from Exim or Sieve filter files, not from traditional forward files. This option is forced to be true if **one\_time** is set.

<b>forbid_filter_run</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------------	----------------------	----------------------	-----------------------

If this option is true, string expansions in Exim filter files are not allowed to make use of **run** items.

<b>forbid_include</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	----------------------	----------------------	-----------------------

If this option is true, items of the form

```
:include:<path name>
```

are not permitted in non-filter redirection lists.

<b>forbid_pipe</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	----------------------	----------------------	-----------------------

If this option is true, this router may not generate a new address which specifies delivery to a pipe, either from an Exim filter or from a conventional forward file. This option is forced to be true if **one\_time** is set.

<b>forbid_sieve_filter</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------------	----------------------	----------------------	-----------------------

If this option is set true, only Exim filters are permitted when **allow\_filter** is true.

<b>forbid_smtp_code</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------	----------------------	----------------------	-----------------------

If this option is set true, any SMTP error codes that are present at the start of messages specified for **:defer:** or **:fail:** are quietly ignored, and the default codes (451 and 550, respectively) are always used.

<b>hide_child_in_errmsg</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------------	----------------------	----------------------	-----------------------

If this option is true, it prevents Exim from quoting a child address if it generates a bounce or delay message for it. Instead it says “an address generated from *<the top level address>*”. Of course, this applies only to bounces generated locally. If a message is forwarded to another host, *its* bounce may well quote the generated address.

<b>ignore_eaccess</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	----------------------	----------------------	-----------------------

If this option is set and an attempt to open a redirection file yields the EACCESS error (permission denied), the *redirect* router behaves as if the file did not exist.

<b>ignore_enotdir</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	----------------------	----------------------	-----------------------

If this option is set and an attempt to open a redirection file yields the ENOTDIR error (something on the path is not a directory), the *redirect* router behaves as if the file did not exist.

Setting **ignore\_enotdir** has another effect as well: When a *redirect* router that has the **file** option set discovers that the file does not exist (the ENOENT error), it tries to *stat()* the parent directory, as a check against unmounted NFS directories. If the parent can not be stattd, delivery is deferred. However, it seems wrong to do this check when **ignore\_enotdir** is set, because that option tells Exim to ignore “something on the path is not a directory” (the ENOTDIR error). This is a confusing area, because it seems that some operating systems give ENOENT where others give ENOTDIR.

<b>include_directory</b>	Use: <i>redirect</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------------------	----------------------	---------------------	-----------------------

If this option is set, the path names of any *:include:* items in a redirection list must start with this directory.

<b>modemask</b>	Use: <i>redirect</i>	Type: <i>octal integer</i>	Default: <i>022</i>
-----------------	----------------------	----------------------------	---------------------

This specifies mode bits which must not be set for a file specified by the **file** option. If any of the forbidden bits are set, delivery is deferred.

<b>one_time</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------	----------------------	----------------------	-----------------------

Sometimes the fact that Exim re-evaluates aliases and reprocesses redirection files each time it tries to deliver a message causes a problem when one or more of the generated addresses fails be delivered at the first attempt. The problem is not one of duplicate delivery – Exim is clever enough to handle that – but of what happens when the redirection list changes during the time that the message is on Exim’s queue. This is particularly true in the case of mailing lists, where new subscribers might receive copies of messages that were posted before they subscribed.

If **one\_time** is set and any addresses generated by the router fail to deliver at the first attempt, the failing addresses are added to the message as “top level” addresses, and the parent address that generated them is marked “delivered”. Thus, redirection does not happen again at the next delivery attempt.

**Warning 1:** Any header line addition or removal that is specified by this router would be lost if delivery did not succeed at the first attempt. For this reason, the **headers\_add** and **headers\_remove** generic options are not permitted when **one\_time** is set.

**Warning 2:** To ensure that the router generates only addresses (as opposed to pipe or file deliveries or auto-replies) **forbid\_file**, **forbid\_pipe**, and **forbid\_filter\_reply** are forced to be true when **one\_time** is set.

**Warning 3:** The **unseen** generic router option may not be set with **one\_time**.

The original top-level address is remembered with each of the generated addresses, and is output in any log messages. However, any intermediate parent addresses are not recorded. This makes a difference to the log only if **all\_parents** log selector is set. It is expected that **one\_time** will typically be used for mailing lists, where there is normally just one level of expansion.

<b>owners</b>	Use: <i>redirect</i>	Type: <i>string list</i>	Default: <i>unset</i>
---------------	----------------------	--------------------------	-----------------------

This specifies a list of permitted owners for the file specified by **file**. This list is in addition to the local user when **check\_local\_user** is set. See **check\_owner** above.

<b>owngroups</b>	Use: <i>redirect</i>	Type: <i>string list</i>	Default: <i>unset</i>
------------------	----------------------	--------------------------	-----------------------

This specifies a list of permitted groups for the file specified by **file**. The list is in addition to the local user's primary group when **check\_local\_user** is set. See **check\_group** above.

<b>pipe_transport</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
-----------------------	----------------------	----------------------	-----------------------

A *redirect* router sets up a direct delivery to a pipe when a string starting with a vertical bar character is specified as a new “address”. The transport used is specified by this option, which, after expansion, must be the name of a configured transport. This should normally be a *pipe* transport. When the transport is run, the pipe command is in *\$address\_pipe*.

<b>qualify_domain</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
-----------------------	----------------------	----------------------	-----------------------

If this option is set, and an unqualified address (one without a domain) is generated, and that address would normally be qualified by the global setting in **qualify\_recipient**, it is instead qualified with the domain specified by expanding this string. If the expansion fails, the router declines. If you want to revert to the default, you can have the expansion generate *\$qualify\_recipient*.

This option applies to all unqualified addresses generated by Exim filters, but for traditional *.forward* files, it applies only to addresses that are not preceded by a backslash. Sieve filters cannot generate unqualified addresses.

<b>qualify_preserve_domain</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------------------	----------------------	----------------------	-----------------------

If this option is set, the router's local **qualify\_domain** option must not be set (a configuration error occurs if it is). If an unqualified address (one without a domain) is generated, it is qualified with the domain of the parent address (the immediately preceding ancestor) instead of the global **qualify\_recipient** value. In the case of a traditional *.forward* file, this applies whether or not the address is preceded by a backslash.

<b>repeat_use</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>true</i>
-------------------	----------------------	----------------------	----------------------

If this option is set false, the router is skipped for a child address that has any ancestor that was routed by this router. This test happens before any of the other preconditions are tested. Exim's default anti-looping rules skip only when the ancestor is the same as the current address. See also **check\_ancestor** above and the generic **redirect\_router** option.

<b>reply_transport</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
------------------------	----------------------	----------------------	-----------------------

A *redirect* router sets up an automatic reply when a **mail** or **vacation** command is used in a filter file. The transport used is specified by this option, which, after expansion, must be the name of a configured transport. This should normally be an *autoreply* transport. Other transports are unlikely to do anything sensible or useful.

<b>rewrite</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>true</i>
----------------	----------------------	----------------------	----------------------

If this option is set false, addresses generated by the router are not subject to address rewriting. Otherwise, they are treated like new addresses and are rewritten according to the global rewriting rules.

<b>sieve_subaddress</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	----------------------	----------------------	-----------------------

The value of this option is passed to a Sieve filter to specify the :subaddress part of an address.

<b>sieve_useraddress</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------------	----------------------	----------------------	-----------------------

The value of this option is passed to a Sieve filter to specify the :user part of an address. However, if it is unset, the entire original local part (including any prefix or suffix) is used for :user.

<b>sieve_vacation_directory</b>	Use: <i>redirect</i>	Type: <i>string†</i>	Default: <i>unset</i>
---------------------------------	----------------------	----------------------	-----------------------

To enable the “vacation” extension for Sieve filters, you must set **sieve\_vacation\_directory** to the directory where vacation databases are held (do not put anything else in that directory), and ensure that the **reply\_transport** option refers to an *autoreply* transport. Each user needs their own directory; Exim will create it if necessary.

<b>skip_syntax_errors</b>	Use: <i>redirect</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	----------------------	----------------------	-----------------------

If **skip\_syntax\_errors** is set, syntactically malformed addresses in non-filter redirection data are skipped, and each failing address is logged. If **syntax\_errors\_to** is set, a message is sent to the address it defines, giving details of the failures. If **syntax\_errors\_text** is set, its contents are expanded and placed at the head of the error message generated by **syntax\_errors\_to**. Usually it is appropriate to set **syntax\_errors\_to** to be the same address as the generic **errors\_to** option. The **skip\_syntax\_errors** option is often used when handling mailing lists.

If all the addresses in a redirection list are skipped because of syntax errors, the router declines to handle the original address, and it is passed to the following routers.

If **skip\_syntax\_errors** is set when an Exim filter is interpreted, any syntax error in the filter causes filtering to be abandoned without any action being taken. The incident is logged, and the router declines to handle the address, so it is passed to the following routers.

Syntax errors in a Sieve filter file cause the “keep” action to occur. This action is specified by RFC 3028. The values of **skip\_syntax\_errors**, **syntax\_errors\_to**, and **syntax\_errors\_text** are not used.

**skip\_syntax\_errors** can be used to specify that errors in users’ forward lists or filter files should not prevent delivery. The **syntax\_errors\_to** option, used with an address that does not get redirected, can be used to notify users of these errors, by means of a router like this:

```

userforward:
  driver = redirect
  allow_filter
  check_local_user
  file = $home/.forward
  file_transport = address_file
  pipe_transport = address_pipe
  reply_transport = address_reply
  no_verify
  skip_syntax_errors
  syntax_errors_to = real-$local_part@$domain
  syntax_errors_text = \
    This is an automatically generated message. An error has\n\
    been found in your .forward file. Details of the error are\n\
    reported below. While this error persists, you will receive\n\
    a copy of this message for every message that is addressed\n\
    to you. If your .forward file is a filter file, or if it is\n\
    a non-filter file containing no valid forwarding addresses,\n\
    a copy of each incoming message will be put in your normal\n\
    mailbox. If a non-filter file contains at least one valid\n\
    forwarding address, forwarding to the valid addresses will\n\
    happen, and those will be the only deliveries that occur.
```

You also need a router to ensure that local addresses that are prefixed by `real-` are recognized, but not forwarded or filtered. For example, you could put this immediately before the *userforward* router:

```
real_localuser:
  driver = accept
  check_local_user
  local_part_prefix = real-
  transport = local_delivery
```

For security, it would probably be a good idea to restrict the use of this router to locally-generated messages, using a condition such as this:

```
condition = ${if match {$sender_host_address}\
                      {\N^(|127\.0\.0\.1)$\N}}
```

<b>syntax_errors_text</b>	Use: <i>redirect</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
---------------------------	----------------------	----------------------------------	-----------------------

See **skip\_syntax\_errors** above.

<b>syntax_errors_to</b>	Use: <i>redirect</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------------	----------------------	---------------------	-----------------------

See **skip\_syntax\_errors** above.



## 23. Environment for running local transports

Local transports handle deliveries to files and pipes. (The *autoreply* transport can be thought of as similar to a pipe.) Exim always runs transports in subprocesses, under specified uids and gids. Typical deliveries to local mailboxes run under the uid and gid of the local user.

Exim also sets a specific current directory while running the transport; for some transports a home directory setting is also relevant. The *pipe* transport is the only one that sets up environment variables; see section 29.4 for details.

The values used for the uid, gid, and the directories may come from several different places. In many cases, the router that handles the address associates settings with that address as a result of its **check\_local\_user**, **group**, or **user** options. However, values may also be given in the transport's own configuration, and these override anything that comes from the router.

### 23.1 Concurrent deliveries

If two different messages for the same local recipient arrive more or less simultaneously, the two delivery processes are likely to run concurrently. When the *appendfile* transport is used to write to a file, Exim applies locking rules to stop concurrent processes from writing to the same file at the same time.

However, when you use a *pipe* transport, it is up to you to arrange any locking that is needed. Here is a silly example:

```
my_transport:
  driver = pipe
  command = /bin/sh -c 'cat >>/some/file'
```

This is supposed to write the message at the end of the file. However, if two messages arrive at the same time, the file will be scrambled. You can use the **exim\_lock** utility program (see section 53.15) to lock a file using the same algorithm that Exim itself uses.

### 23.2 Uids and gids

All transports have the options **group** and **user**. If **group** is set, it overrides any group that the router set in the address, even if **user** is not set for the transport. This makes it possible, for example, to run local mail delivery under the uid of the recipient (set by the router), but in a special group (set by the transport). For example:

```
# Routers ...
# User/group are set by check_local_user in this router
local_users:
  driver = accept
  check_local_user
  transport = group_delivery

# Transports ...
# This transport overrides the group
group_delivery:
  driver = appendfile
  file = /var/spool/mail/$local_part
  group = mail
```

If **user** is set for a transport, its value overrides what is set in the address by the router. If **user** is non-numeric and **group** is not set, the gid associated with the user is used. If **user** is numeric, **group** must be set.

When the uid is taken from the transport's configuration, the *initgroups()* function is called for the groups associated with that uid if the **initgroups** option is set for the transport. When the uid is not

specified by the transport, but is associated with the address by a router, the option for calling *initgroups()* is taken from the router configuration.

The *pipe* transport contains the special option **pipe\_as\_creator**. If this is set and **user** is not set, the uid of the process that called Exim to receive the message is used, and if **group** is not set, the corresponding original gid is also used.

This is the detailed preference order for obtaining a gid; the first of the following that is set is used:

- A **group** setting of the transport;
- A **group** setting of the router;
- A gid associated with a user setting of the router, either as a result of **check\_local\_user** or an explicit non-numeric **user** setting;
- The group associated with a non-numeric **user** setting of the transport;
- In a *pipe* transport, the creator's gid if **deliver\_as\_creator** is set and the uid is the creator's uid;
- The Exim gid if the Exim uid is being used as a default.

If, for example, the user is specified numerically on the router and there are no group settings, no gid is available. In this situation, an error occurs. This is different for the uid, for which there always is an ultimate default. The first of the following that is set is used:

- A **user** setting of the transport;
- In a *pipe* transport, the creator's uid if **deliver\_as\_creator** is set;
- A **user** setting of the router;
- A **check\_local\_user** setting of the router;
- The Exim uid.

Of course, an error will still occur if the uid that is chosen is on the **never\_users** list.

## 23.3 Current and home directories

Routers may set current and home directories for local transports by means of the **transport\_current\_directory** and **transport\_home\_directory** options. However, if the transport's **current\_directory** or **home\_directory** options are set, they override the router's values. In detail, the home directory for a local transport is taken from the first of these values that is set:

- The **home\_directory** option on the transport;
- The **transport\_home\_directory** option on the router;
- The password data if **check\_local\_user** is set on the router;
- The **router\_home\_directory** option on the router.

The current directory is taken from the first of these values that is set:

- The **current\_directory** option on the transport;
- The **transport\_current\_directory** option on the router.

If neither the router nor the transport sets a current directory, Exim uses the value of the home directory, if it is set. Otherwise it sets the current directory to / before running a local transport.

## 23.4 Expansion variables derived from the address

Normally a local delivery is handling a single address, and in that case the variables such as *\$domain* and *\$local\_part* are set during local deliveries. However, in some circumstances more than one address may be handled at once (for example, while writing batch SMTP for onward transmission by some other means). In this case, the variables associated with the local part are never set, *\$domain* is set only if all the addresses have the same domain, and *\$original\_domain* is never set.

## 24. Generic options for transports

The following generic options apply to all transports:

<b>body_only</b>	Use: <i>transports</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------	------------------------	----------------------	-----------------------

If this option is set, the message's headers are not transported. It is mutually exclusive with **headers\_only**. If it is used with the *appendfile* or *pipe* transports, the settings of **message\_prefix** and **message\_suffix** should be checked, because this option does not automatically suppress them.

<b>current_directory</b>	Use: <i>transports</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
--------------------------	------------------------	----------------------------------	-----------------------

This specifies the current directory that is to be set while running the transport, overriding any value that may have been set by the router. If the expansion fails for any reason, including forced failure, an error is logged, and delivery is deferred.

<b>disable_logging</b>	Use: <i>transports</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------------	----------------------	-----------------------

If this option is set true, nothing is logged for any deliveries by the transport or for any transport errors. You should not set this option unless you really, really know what you are doing.

<b>debug_print</b>	Use: <i>transports</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
--------------------	------------------------	----------------------------------	-----------------------

If this option is set and debugging is enabled (see the **-d** command line option), the string is expanded and included in the debugging output when the transport is run. If expansion of the string fails, the error message is written to the debugging output, and Exim carries on processing. This facility is provided to help with checking out the values of variables and so on when debugging driver configurations. For example, if a **headers\_add** option is not working properly, **debug\_print** could be used to output the variables it references. A newline is added to the text if it does not end with one. The variables *\$transport\_name* and *\$router\_name* contain the name of the transport and the router that called it.

<b>delivery_date_add</b>	Use: <i>transports</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------------	------------------------	----------------------	-----------------------

If this option is true, a *Delivery-date:* header is added to the message. This gives the actual time the delivery was made. As this is not a standard header, Exim has a configuration option (**delivery\_date\_remove**) which requests its removal from incoming messages, so that delivered messages can safely be resent to other recipients.

<b>driver</b>	Use: <i>transports</i>	Type: <i>string</i>	Default: <i>unset</i>
---------------	------------------------	---------------------	-----------------------

This specifies which of the available transport drivers is to be used. There is no default, and this option must be set for every transport.

<b>envelope_to_add</b>	Use: <i>transports</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------------	----------------------	-----------------------

If this option is true, an *Envelope-to:* header is added to the message. This gives the original address(es) in the incoming envelope that caused this delivery to happen. More than one address may be present if the transport is configured to handle several addresses at once, or if more than one original address was redirected to the same final address. As this is not a standard header, Exim has a configuration option (**envelope\_to\_remove**) which requests its removal from incoming messages, so that delivered messages can safely be resent to other recipients.

<b>event_action</b>	Use: <i>transports</i>	Type: <i>string</i> †	Default: <i>unset</i>
---------------------	------------------------	-----------------------	-----------------------

This option declares a string to be expanded for Exim's events mechanism. For details see chapter 60.

<b>group</b>	Use: <i>transports</i>	Type: <i>string</i> †	Default: <i>Exim group</i>
--------------	------------------------	-----------------------	----------------------------

This option specifies a gid for running the transport process, overriding any value that the router supplies, and also overriding any value associated with **user** (see below).

<b>headers_add</b>	Use: <i>transports</i>	Type: <i>list</i> †	Default: <i>unset</i>
--------------------	------------------------	---------------------	-----------------------

This option specifies a list of text headers, newline-separated (by default, changeable in the usual way), which are (separately) expanded and added to the header portion of a message as it is transported, as described in section 47.17. Additional header lines can also be specified by routers. If the result of the expansion is an empty string, or if the expansion is forced to fail, no action is taken. Other expansion failures are treated as errors and cause the delivery to be deferred.

Unlike most options, **headers\_add** can be specified multiple times for a transport; all listed headers are added.

<b>headers_only</b>	Use: <i>transports</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------	------------------------	----------------------	-----------------------

If this option is set, the message's body is not transported. It is mutually exclusive with **body\_only**. If it is used with the *appendfile* or *pipe* transports, the settings of **message\_prefix** and **message\_suffix** should be checked, since this option does not automatically suppress them.

<b>headers_remove</b>	Use: <i>transports</i>	Type: <i>list</i> †	Default: <i>unset</i>
-----------------------	------------------------	---------------------	-----------------------

This option specifies a list of header names, colon-separated (by default, changeable in the usual way); these headers are omitted from the message as it is transported, as described in section 47.17. Header removal can also be specified by routers. Each list item is separately expanded. If the result of the expansion is an empty string, or if the expansion is forced to fail, no action is taken. Other expansion failures are treated as errors and cause the delivery to be deferred.

Unlike most options, **headers\_remove** can be specified multiple times for a transport; all listed headers are removed.

**Warning:** Because of the separate expansion of the list items, items that contain a list separator must have it doubled. To avoid this, change the list separator (6.21).

<b>headers_rewrite</b>	Use: <i>transports</i>	Type: <i>string</i>	Default: <i>unset</i>
------------------------	------------------------	---------------------	-----------------------

This option allows addresses in header lines to be rewritten at transport time, that is, as the message is being copied to its destination. The contents of the option are a colon-separated list of rewriting rules. Each rule is in exactly the same form as one of the general rewriting rules that are applied when a message is received. These are described in chapter 31. For example,

```
headers_rewrite = a@b c@d f : \
                  x@y w@z
```

changes *a@b* into *c@d* in *From:* header lines, and *x@y* into *w@z* in all address-bearing header lines. The rules are applied to the header lines just before they are written out at transport time, so they affect only those copies of the message that pass through the transport. However, only the message's original header lines, and any that were added by a system filter, are rewritten. If a router or transport adds header lines, they are not affected by this option. These rewriting rules are *not* applied to the envelope. You can change the return path using **return\_path**, but you cannot change envelope recipients at this time.

<b>home_directory</b>	Use: <i>transports</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-----------------------	------------------------	----------------------------------	-----------------------

This option specifies a home directory setting for a local transport, overriding any value that may be set by the router. The home directory is placed in *\$home* while expanding the transport's private options. It is also used as the current directory if no current directory is set by the **current\_directory** option on the transport or the **transport\_current\_directory** option on the router. If the expansion fails for any reason, including forced failure, an error is logged, and delivery is deferred.

<b>initgroups</b>	Use: <i>transports</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------	------------------------	----------------------	-----------------------

If this option is true and the uid for the delivery process is provided by the transport, the *initgroups()* function is called when running the transport to ensure that any additional groups associated with the uid are set up.

<b>max_parallel</b>	Use: <i>transports</i>	Type: <i>integer</i> <sup>†</sup>	Default: <i>unset</i>
---------------------	------------------------	-----------------------------------	-----------------------

If this option is set and expands to an integer greater than zero it limits the number of concurrent runs of the transport. The control does not apply to shadow transports.

Exim implements this control by means of a hints database in which a record is incremented whenever a transport process is being created. The record is decremented and possibly removed when the process terminates. Obviously there is scope for records to get left lying around if there is a system or program crash. To guard against this, Exim ignores any records that are more than six hours old.

If you use this option, you should also arrange to delete the relevant hints database whenever your system reboots. The names of the files start with *misc* and they are kept in the *spool/db* directory. There may be one or two files, depending on the type of DBM in use. The same files are used for ETRN and smtp transport serialization.

<b>message_size_limit</b>	Use: <i>transports</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>0</i>
---------------------------	------------------------	----------------------------------	-------------------

This option controls the size of messages passed through the transport. It is expanded before use; the result of the expansion must be a sequence of decimal digits, optionally followed by K or M. If the expansion fails for any reason, including forced failure, or if the result is not of the required form, delivery is deferred. If the value is greater than zero and the size of a message exceeds this limit, the address is failed. If there is any chance that the resulting bounce message could be routed to the same transport, you should ensure that **return\_size\_limit** is less than the transport's **message\_size\_limit**, as otherwise the bounce message will fail to get delivered.

<b>rcpt_include_affixes</b>	Use: <i>transports</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------------	------------------------	----------------------	-----------------------

When this option is false (the default), and an address that has had any affixes (prefixes or suffixes) removed from the local part is delivered by any form of SMTP or LMTP, the affixes are not included. For example, if a router that contains

```
local_part_prefix = *-
```

routes the address *abc-xyz@some.domain* to an SMTP transport, the envelope is delivered with

```
RCPT TO:<xyz@some.domain>
```

This is also the case when an ACL-time callout is being used to verify a recipient address. However, if **rcpt\_include\_affixes** is set true, the whole local part is included in the RCPT command. This option applies to BSMTP deliveries by the *appendfile* and *pipe* transports as well as to the *lmtp* and *smtp* transports.

<b>retry_use_local_part</b>	Use: <i>transports</i>	Type: <i>boolean</i>	Default: <i>see below</i>
-----------------------------	------------------------	----------------------	---------------------------

When a delivery suffers a temporary failure, a retry record is created in Exim's hints database. For remote deliveries, the key for the retry record is based on the name and/or IP address of the failing remote host. For local deliveries, the key is normally the entire address, including both the local part and the domain. This is suitable for most common cases of local delivery temporary failure – for example, exceeding a mailbox quota should delay only deliveries to that mailbox, not to the whole domain.

However, in some special cases you may want to treat a temporary local delivery as a failure associated with the domain, and not with a particular local part. (For example, if you are storing all mail for some domain in files.) You can do this by setting **retry\_use\_local\_part** false.

For all the local transports, its default value is true. For remote transports, the default value is false for tidiness, but changing the value has no effect on a remote transport in the current implementation.

<b>return_path</b>	Use: <i>transports</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------	------------------------	----------------------	-----------------------

If this option is set, the string is expanded at transport time and replaces the existing return path (envelope sender) value in the copy of the message that is being delivered. An empty return path is permitted. This feature is designed for remote deliveries, where the value of this option is used in the SMTP MAIL command. If you set **return\_path** for a local transport, the only effect is to change the address that is placed in the *Return-path:* header line, if one is added to the message (see the next option).

**Note:** A changed return path is not logged unless you add **return\_path\_on\_delivery** to the log selector.

The expansion can refer to the existing value via *\$return\_path*. This is either the message's envelope sender, or an address set by the **errors\_to** option on a router. If the expansion is forced to fail, no replacement occurs; if it fails for another reason, delivery is deferred. This option can be used to support VERP (Variable Envelope Return Paths) – see section 50.6.

**Note:** If a delivery error is detected locally, including the case when a remote server rejects a message at SMTP time, the bounce message is not sent to the value of this option. It is sent to the previously set errors address. This defaults to the incoming sender address, but can be changed by setting **errors\_to** in a router.

<b>return_path_add</b>	Use: <i>transports</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------------	----------------------	-----------------------

If this option is true, a *Return-path:* header is added to the message. Although the return path is normally available in the prefix line of BSD mailboxes, this is commonly not displayed by MUAs, and so the user does not have easy access to it.

RFC 2821 states that the *Return-path:* header is added to a message “when the delivery SMTP server makes the final delivery”. This implies that this header should not be present in incoming messages. Exim has a configuration option, **return\_path\_remove**, which requests removal of this header from incoming messages, so that delivered messages can safely be resent to other recipients.

<b>shadow_condition</b>	Use: <i>transports</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	------------------------	----------------------	-----------------------

See **shadow\_transport** below.

<b>shadow_transport</b>	Use: <i>transports</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------------	------------------------	---------------------	-----------------------

A local transport may set the **shadow\_transport** option to the name of another local transport. Shadow remote transports are not supported.

Whenever a delivery to the main transport succeeds, and either **shadow\_condition** is unset, or its expansion does not result in the empty string or one of the strings “0” or “no” or “false”, the message is also passed to the shadow transport, with the same delivery address or addresses. If expansion fails, no action is taken except that non-forced expansion failures cause a log line to be written.

The result of the shadow transport is discarded and does not affect the subsequent processing of the message. Only a single level of shadowing is provided; the **shadow\_transport** option is ignored on any transport when it is running as a shadow. Options concerned with output from pipes are also ignored. The log line for the successful delivery has an item added on the end, of the form

ST=<shadow transport name>

If the shadow transport did not succeed, the error message is put in parentheses afterwards. Shadow transports can be used for a number of different purposes, including keeping more detailed log information than Exim normally provides, and implementing automatic acknowledgment policies based on message headers that some sites insist on.

<b>transport_filter</b>	Use: <i>transports</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-------------------------	------------------------	----------------------------------	-----------------------

This option sets up a filtering (in the Unix shell sense) process for messages at transport time. It should not be confused with mail filtering as set up by individual users or via a system filter. If unset, or expanding to an empty string, no filtering is done.

When the message is about to be written out, the command specified by **transport\_filter** is started up in a separate, parallel process, and the entire message, including the header lines, is passed to it on its standard input (this in fact is done from a third process, to avoid deadlock). The command must be specified as an absolute path.

The lines of the message that are written to the transport filter are terminated by newline (“\n”). The message is passed to the filter before any SMTP-specific processing, such as turning “\n” into “\r\n” and escaping lines beginning with a dot, and also before any processing implied by the settings of **check\_string** and **escape\_string** in the *appendfile* or *pipe* transports.

The standard error for the filter process is set to the same destination as its standard output; this is read and written to the message’s ultimate destination. The process that writes the message to the filter, the filter itself, and the original process that reads the result and delivers it are all run in parallel, like a shell pipeline.

The filter can perform any transformations it likes, but of course should take care not to break RFC 2822 syntax. Exim does not check the result, except to test for a final newline when SMTP is in use. All messages transmitted over SMTP must end with a newline, so Exim supplies one if it is missing.

A transport filter can be used to provide content-scanning on a per-user basis at delivery time if the only required effect of the scan is to modify the message. For example, a content scan could insert a new header line containing a spam score. This could be interpreted by a filter in the user’s MUA. It is not possible to discard a message at this stage.

A problem might arise if the filter increases the size of a message that is being sent down an SMTP connection. If the receiving SMTP server has indicated support for the SIZE parameter, Exim will have sent the size of the message at the start of the SMTP session. If what is actually sent is substantially more, the server might reject the message. This can be worked round by setting the **size\_addition** option on the *smtp* transport, either to allow for additions to the message, or to disable the use of SIZE altogether.

The value of the **transport\_filter** option is the command string for starting the filter, which is run directly from Exim, not under a shell. The string is parsed by Exim in the same way as a command string for the *pipe* transport: Exim breaks it up into arguments and then expands each argument separately (see section 29.3). Any kind of expansion failure causes delivery to be deferred. The special argument *\$pipe\_addresses* is replaced by a number of arguments, one for each address that applies to this delivery. (This isn’t an ideal name for this feature here, but as it was already implemented for the *pipe* transport, it seemed sensible not to change it.)

The expansion variables *\$host* and *\$host\_address* are available when the transport is a remote one. They contain the name and IP address of the host to which the message is being sent. For example:

```
transport_filter = /some/directory/transport-filter.pl \  
    $host $host_address $sender_address $pipe_addresses
```

Two problems arise if you want to use more complicated expansion items to generate transport filter commands, both of which due to the fact that the command is split up *before* expansion.

- If an expansion item contains white space, you must quote it, so that it is all part of the same command item. If the entire option setting is one such expansion item, you have to take care what kind of quoting you use. For example:

```
transport_filter = '/bin/cmd${if eq{$host}{a.b.c}{1}{2}}'
```

This runs the command */bin/cmd1* if the host name is *a.b.c*, and */bin/cmd2* otherwise. If double quotes had been used, they would have been stripped by Exim when it read the option's value. When the value is used, if the single quotes were missing, the line would be split into two items, */bin/cmd\${if* and *eq{\$host}{a.b.c}{1}{2}}*, and an error would occur when Exim tried to expand the first one.

- Except for the special case of *\$pipe\_addresses* that is mentioned above, an expansion cannot generate multiple arguments, or a command name followed by arguments. Consider this example:

```
transport_filter = ${lookup{$host}lsearch{/a/file}\  
    {$value}{/bin/cat}}
```

The result of the lookup is interpreted as the name of the command, even if it contains white space. The simplest way round this is to use a shell:

```
transport_filter = /bin/sh -c ${lookup{$host}lsearch{/a/file}\  
    {$value}{/bin/cat}}
```

The filter process is run under the same uid and gid as the normal delivery. For remote deliveries this is the Exim uid/gid by default. The command should normally yield a zero return code. Transport filters are not supposed to fail. A non-zero code is taken to mean that the transport filter encountered some serious problem. Delivery of the message is deferred; the message remains on the queue and is tried again later. It is not possible to cause a message to be bounced from a transport filter.

If a transport filter is set on an autoreply transport, the original message is passed through the filter as it is being copied into the newly generated message, which happens if the **return\_message** option is set.

<b>transport_filter_timeout</b>	Use: <i>transports</i>	Type: <i>time</i>	Default: <i>5m</i>
---------------------------------	------------------------	-------------------	--------------------

When Exim is reading the output of a transport filter, it applies a timeout that can be set by this option. Exceeding the timeout is normally treated as a temporary delivery failure. However, if a transport filter is used with a *pipe* transport, a timeout in the transport filter is treated in the same way as a timeout in the pipe command itself. By default, a timeout is a hard error, but if the *pipe* transport's **timeout\_defer** option is set true, it becomes a temporary error.

<b>user</b>	Use: <i>transports</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>Exim user</i>
-------------	------------------------	----------------------------------	---------------------------

This option specifies the user under whose uid the delivery process is to be run, overriding any uid that may have been set by the router. If the user is given as a name, the uid is looked up from the password data, and the associated group is taken as the value of the gid to be used if the **group** option is not set.

For deliveries that use local transports, a user and group are normally specified explicitly or implicitly (for example, as a result of **check\_local\_user**) by the router or transport.



For remote transports, you should leave this option unset unless you really are sure you know what you are doing. When a remote transport is running, it needs to be able to access Exim's hints databases, because each host may have its own retry data.

## 25. Address batching in local transports

The only remote transport (*smtp*) is normally configured to handle more than one address at a time, so that when several addresses are routed to the same remote host, just one copy of the message is sent. Local transports, however, normally handle one address at a time. That is, a separate instance of the transport is run for each address that is routed to the transport. A separate copy of the message is delivered each time.

In special cases, it may be desirable to handle several addresses at once in a local transport, for example:

- In an *appendfile* transport, when storing messages in files for later delivery by some other means, a single copy of the message with multiple recipients saves space.
- In an *lmtp* transport, when delivering over “local SMTP” to some process, a single copy saves time, and is the normal way LMTP is expected to work.
- In a *pipe* transport, when passing the message to a scanner program or to some other delivery mechanism such as UUCP, multiple recipients may be acceptable.

These three local transports all have the same options for controlling multiple (“batched”) deliveries, namely **batch\_max** and **batch\_id**. To save repeating the information for each transport, these options are described here.

The **batch\_max** option specifies the maximum number of addresses that can be delivered together in a single run of the transport. Its default value is one (no batching). When more than one address is routed to a transport that has a **batch\_max** value greater than one, the addresses are delivered in a batch (that is, in a single run of the transport with multiple recipients), subject to certain conditions:

- If any of the transport’s options contain a reference to *\$local\_part*, no batching is possible.
- If any of the transport’s options contain a reference to *\$domain*, only addresses with the same domain are batched.
- If **batch\_id** is set, it is expanded for each address, and only those addresses with the same expanded value are batched. This allows you to specify customized batching conditions. Failure of the expansion for any reason, including forced failure, disables batching, but it does not stop the delivery from taking place.
- Batched addresses must also have the same errors address (where to send delivery errors), the same header additions and removals, the same user and group for the transport, and if a host list is present, the first host must be the same.

In the case of the *appendfile* and *pipe* transports, batching applies both when the file or pipe command is specified in the transport, and when it is specified by a *redirect* router, but all the batched addresses must of course be routed to the same file or pipe command. These two transports have an option called **use\_bsmtplib**, which causes them to deliver the message in “batched SMTP” format, with the envelope represented as SMTP commands. The **check\_string** and **escape\_string** options are forced to the values

```
check_string = "."
escape_string = ".."
```

when batched SMTP is in use. A full description of the batch SMTP mechanism is given in section 48.10. The *lmtp* transport does not have a **use\_bsmtplib** option, because it always delivers using the SMTP protocol.

If the generic **envelope\_to\_add** option is set for a batching transport, the *Envelope-to:* header that is added to the message contains all the addresses that are being processed together. If you are using a batching *appendfile* transport without **use\_bsmtplib**, the only way to preserve the recipient addresses is to set the **envelope\_to\_add** option.

If you are using a *pipe* transport without BSMTP, and setting the transport’s **command** option, you can include *\$pipe\_addresses* as part of the command. This is not a true variable; it is a bit of magic

that causes each of the recipient addresses to be inserted into the command as a separate argument. This provides a way of accessing all the addresses that are being delivered in the batch. **Note:** This is not possible for pipe commands that are specified by a *redirect* router.

## 26. The appendfile transport

The *appendfile* transport delivers a message by appending it to an existing file, or by creating an entirely new file in a specified directory. Single files to which messages are appended can be in the traditional Unix mailbox format, or optionally in the MBX format supported by the Pine MUA and University of Washington IMAP daemon, *inter alia*. When each message is being delivered as a separate file, “maildir” format can optionally be used to give added protection against failures that happen part-way through the delivery. A third form of separate-file delivery known as “mailstore” is also supported. For all file formats, Exim attempts to create as many levels of directory as necessary, provided that **create\_directory** is set.

The code for the optional formats is not included in the Exim binary by default. It is necessary to set **SUPPORT\_MBX**, **SUPPORT\_MAILDIR** and/or **SUPPORT\_MAILSTORE** in *Local/Makefile* to have the appropriate code included.

Exim recognizes system quota errors, and generates an appropriate message. Exim also supports its own quota control within the transport, for use when the system facility is unavailable or cannot be used for some reason.

If there is an error while appending to a file (for example, quota exceeded or partition filled), Exim attempts to reset the file’s length and last modification time back to what they were before. If there is an error while creating an entirely new file, the new file is removed.

Before appending to a file, a number of security checks are made, and the file is locked. A detailed description is given below, after the list of private options.

The *appendfile* transport is most commonly used for local deliveries to users’ mailboxes. However, it can also be used as a pseudo-remote transport for putting messages into files for remote delivery by some means other than Exim. “Batch SMTP” format is often used in this case (see the **use\_bsmtip** option).

### 26.1 The file and directory options

The **file** option specifies a single file, to which the message is appended; the **directory** option specifies a directory, in which a new file containing the message is created. Only one of these two options can be set, and for normal deliveries to mailboxes, one of them *must* be set.

However, *appendfile* is also used for delivering messages to files or directories whose names (or parts of names) are obtained from alias, forwarding, or filtering operations (for example, a **save** command in a user’s Exim filter). When such a transport is running, *\$local\_part* contains the local part that was aliased or forwarded, and *\$address\_file* contains the name (or partial name) of the file or directory generated by the redirection operation. There are two cases:

- If neither **file** nor **directory** is set, the redirection operation must specify an absolute path (one that begins with /). This is the most common case when users with local accounts use filtering to sort mail into different folders. See for example, the *address\_file* transport in the default configuration. If the path ends with a slash, it is assumed to be the name of a directory. A delivery to a directory can also be forced by setting **maildir\_format** or **mailstore\_format**.
- If **file** or **directory** is set for a delivery from a redirection, it is used to determine the file or directory name for the delivery. Normally, the contents of *\$address\_file* are used in some way in the string expansion.

As an example of the second case, consider an environment where users do not have home directories. They may be permitted to use Exim filter commands of the form:

```
save folder23
```

or Sieve filter commands of the form:

```
require "fileinto";  
fileinto "folder23";
```

In this situation, the expansion of **file** or **directory** in the transport must transform the relative path into an appropriate absolute file name. In the case of Sieve filters, the name *inbox* must be handled. It is the name that is used as a result of a “keep” action in the filter. This example shows one way of handling this requirement:

```
file = ${if eq{$address_file}{inbox} \
    {/var/mail/$local_part} \
    ${if eq${substr_0_1:$address_file}{/} \
        {$address_file} \
        {$home/mail/$address_file} \
    }} \
}
```

With this setting of **file**, *inbox* refers to the standard mailbox location, absolute paths are used without change, and other folders are in the *mail* directory within the home directory.

**Note 1:** While processing an Exim filter, a relative path such as *folder23* is turned into an absolute path if a home directory is known to the router. In particular, this is the case if **check\_local\_user** is set. If you want to prevent this happening at routing time, you can set **router\_home\_directory** empty. This forces the router to pass the relative path to the transport.

**Note 2:** An absolute path in *\$address\_file* is not treated specially; the **file** or **directory** option is still used if it is set.

## 26.2 Private options for appendfile

<b>allow_fifo</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------	------------------------	----------------------	-----------------------

Setting this option permits delivery to named pipes (FIFOs) as well as to regular files. If no process is reading the named pipe at delivery time, the delivery is deferred.

<b>allow_symlink</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------------	----------------------	-----------------------

By default, *appendfile* will not deliver if the path name for the file is that of a symbolic link. Setting this option relaxes that constraint, but there are security issues involved in the use of symbolic links. Be sure you know what you are doing if you set this. Details of exactly what this option affects are included in the discussion which follows this list of options.

<b>batch_id</b>	Use: <i>appendfile</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-----------------	------------------------	----------------------------------	-----------------------

See the description of local delivery batching in chapter 25. However, batching is automatically disabled for *appendfile* deliveries that happen as a result of forwarding or aliasing or other redirection directly to a file.

<b>batch_max</b>	Use: <i>appendfile</i>	Type: <i>integer</i>	Default: <i>1</i>
------------------	------------------------	----------------------	-------------------

See the description of local delivery batching in chapter 25.

<b>check_group</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	------------------------	----------------------	-----------------------

When this option is set, the group owner of the file defined by the **file** option is checked to see that it is the same as the group under which the delivery process is running. The default setting is false because the default file mode is 0600, which means that the group is irrelevant.

<b>check_owner</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>true</i>
--------------------	------------------------	----------------------	----------------------

When this option is set, the owner of the file defined by the **file** option is checked to ensure that it is the same as the user under which the delivery process is running.

<b>check_string</b>	Use: <i>appendfile</i>	Type: <i>string</i>	Default: <i>see below</i>
---------------------	------------------------	---------------------	---------------------------

As *appendfile* writes the message, the start of each line is tested for matching **check\_string**, and if it does, the initial matching characters are replaced by the contents of **escape\_string**. The value of **check\_string** is a literal string, not a regular expression, and the case of any letters it contains is significant.

If **use\_bsmtp** is set the values of **check\_string** and **escape\_string** are forced to “.” and “.” respectively, and any settings in the configuration are ignored. Otherwise, they default to “From ” and “>From ” when the **file** option is set, and unset when any of the **directory**, **maildir**, or **mailstore** options are set.

The default settings, along with **message\_prefix** and **message\_suffix**, are suitable for traditional “BSD” mailboxes, where a line beginning with “From ” indicates the start of a new message. All four options need changing if another format is used. For example, to deliver to mailboxes in MMDf format:

```
check_string = "\1\1\1\1\n"
escape_string = "\1\1\1\1 \n"
message_prefix = "\1\1\1\1\n"
message_suffix = "\1\1\1\1\n"
```

<b>create_directory</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>true</i>
-------------------------	------------------------	----------------------	----------------------

When this option is true, Exim attempts to create any missing superior directories for the file that it is about to write. A created directory’s mode is given by the **directory\_mode** option.

The group ownership of a newly created directory is highly dependent on the operating system (and possibly the file system) that is being used. For example, in Solaris, if the parent directory has the setgid bit set, its group is propagated to the child; if not, the currently set group is used. However, in FreeBSD, the parent’s group is always used.

<b>create_file</b>	Use: <i>appendfile</i>	Type: <i>string</i>	Default: <i>anywhere</i>
--------------------	------------------------	---------------------	--------------------------

This option constrains the location of files and directories that are created by this transport. It applies to files defined by the **file** option and directories defined by the **directory** option. In the case of maildir delivery, it applies to the top level directory, not the maildir directories beneath.

The option must be set to one of the words “anywhere”, “inhome”, or “belowhome”. In the second and third cases, a home directory must have been set for the transport. This option is not useful when an explicit file name is given for normal mailbox deliveries. It is intended for the case when file names are generated from users’ *.forward* files. These are usually handled by an *appendfile* transport called **address\_file**. See also **file\_must\_exist**.

<b>directory</b>	Use: <i>appendfile</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
------------------	------------------------	----------------------------------	-----------------------

This option is mutually exclusive with the **file** option, but one of **file** or **directory** must be set, unless the delivery is the direct result of a redirection (see section 26.1).

When **directory** is set, the string is expanded, and the message is delivered into a new file or files in or below the given directory, instead of being appended to a single mailbox file. A number of different formats are provided (see **maildir\_format** and **mailstore\_format**), and see section 26.4 for further details of this form of delivery.

<b>directory_file</b>	Use: <i>appendfile</i>	Type: <i>string†</i>	Default: <i>see below</i>
-----------------------	------------------------	----------------------	---------------------------

When **directory** is set, but neither **maildir\_format** nor **mailstore\_format** is set, *appendfile* delivers each message into a file whose name is obtained by expanding this string. The default value is:

```
q${base62:$tod_epoch}-${inode}
```

This generates a unique name from the current time, in base 62 form, and the inode of the file. The variable *\$inode* is available only when expanding this option.

<b>directory_mode</b>	Use: <i>appendfile</i>	Type: <i>octal integer</i>	Default: <i>0700</i>
-----------------------	------------------------	----------------------------	----------------------

If *appendfile* creates any directories as a result of the **create\_directory** option, their mode is specified by this option.

<b>escape_string</b>	Use: <i>appendfile</i>	Type: <i>string</i>	Default: <i>see description</i>
----------------------	------------------------	---------------------	---------------------------------

See **check\_string** above.

<b>file</b>	Use: <i>appendfile</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------	------------------------	----------------------	-----------------------

This option is mutually exclusive with the **directory** option, but one of **file** or **directory** must be set, unless the delivery is the direct result of a redirection (see section 26.1). The **file** option specifies a single file, to which the message is appended. One or more of **use\_fcntl\_lock**, **use\_flock\_lock**, or **use\_lockfile** must be set with **file**.

If you are using more than one host to deliver over NFS into the same mailboxes, you should always use lock files.

The string value is expanded for each delivery, and must yield an absolute path. The most common settings of this option are variations on one of these examples:

```
file = /var/spool/mail/$local_part
file = /home/$local_part/inbox
file = $home/inbox
```

In the first example, all deliveries are done into the same directory. If Exim is configured to use lock files (see **use\_lockfile** below) it must be able to create a file in the directory, so the “sticky” bit must be turned on for deliveries to be possible, or alternatively the **group** option can be used to run the delivery under a group id which has write access to the directory.

<b>file_format</b>	Use: <i>appendfile</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------------	------------------------	---------------------	-----------------------

This option requests the transport to check the format of an existing file before adding to it. The check consists of matching a specific string at the start of the file. The value of the option consists of an even number of colon-separated strings. The first of each pair is the test string, and the second is the name of a transport. If the transport associated with a matched string is not the current transport, control is passed over to the other transport. For example, suppose the standard *local\_delivery* transport has this added to it:

```
file_format = "From      : local_delivery :\n\\1\\1\\1\\1\\n : local_mmdf_delivery"
```

Mailboxes that begin with “From” are still handled by this transport, but if a mailbox begins with four binary ones followed by a newline, control is passed to a transport called **local\_mmdf\_delivery**, which presumably is configured to do the delivery in MMDF format. If a mailbox does not exist or is empty, it is assumed to match the current transport. If the start of a mailbox doesn’t match any string, or if the transport named for a given string is not defined, delivery is deferred.

<b>file_must_exist</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------------	----------------------	-----------------------

If this option is true, the file specified by the **file** option must exist. A temporary error occurs if it does not, causing delivery to be deferred. If this option is false, the file is created if it does not exist.

<b>lock_fcntl_timeout</b>	Use: <i>appendfile</i>	Type: <i>time</i>	Default: <i>0s</i>
---------------------------	------------------------	-------------------	--------------------

By default, the *appendfile* transport uses non-blocking calls to *fcntl()* when locking an open mailbox file. If the call fails, the delivery process sleeps for **lock\_interval** and tries again, up to **lock\_retries** times. Non-blocking calls are used so that the file is not kept open during the wait for the lock; the reason for this is to make it as safe as possible for deliveries over NFS in the case when processes might be accessing an NFS mailbox without using a lock file. This should not be done, but misunderstandings and hence misconfigurations are not unknown.

On a busy system, however, the performance of a non-blocking lock approach is not as good as using a blocking lock with a timeout. In this case, the waiting is done inside the system call, and Exim's delivery process acquires the lock and can proceed as soon as the previous lock holder releases it.

If **lock\_fcntl\_timeout** is set to a non-zero time, blocking locks, with that timeout, are used. There may still be some retrying: the maximum number of retries is

$$(\text{lock\_retries} * \text{lock\_interval}) / \text{lock\_fcntl\_timeout}$$

rounded up to the next whole number. In other words, the total time during which *appendfile* is trying to get a lock is roughly the same, unless **lock\_fcntl\_timeout** is set very large.

You should consider setting this option if you are getting a lot of delayed local deliveries because of errors of the form

```
failed to lock mailbox /some/file (fcntl)
```

<b>lock_flock_timeout</b>	Use: <i>appendfile</i>	Type: <i>time</i>	Default: <i>0s</i>
---------------------------	------------------------	-------------------	--------------------

This timeout applies to file locking when using *flock()* (see **use\_flock**); the timeout operates in a similar manner to **lock\_fcntl\_timeout**.

<b>lock_interval</b>	Use: <i>appendfile</i>	Type: <i>time</i>	Default: <i>3s</i>
----------------------	------------------------	-------------------	--------------------

This specifies the time to wait between attempts to lock the file. See below for details of locking.

<b>lock_retries</b>	Use: <i>appendfile</i>	Type: <i>integer</i>	Default: <i>10</i>
---------------------	------------------------	----------------------	--------------------

This specifies the maximum number of attempts to lock the file. A value of zero is treated as 1. See below for details of locking.

<b>lockfile_mode</b>	Use: <i>appendfile</i>	Type: <i>octal integer</i>	Default: <i>0600</i>
----------------------	------------------------	----------------------------	----------------------

This specifies the mode of the created lock file, when a lock file is being used (see **use\_lockfile** and **use\_mbx\_lock**).

<b>lockfile_timeout</b>	Use: <i>appendfile</i>	Type: <i>time</i>	Default: <i>30m</i>
-------------------------	------------------------	-------------------	---------------------

When a lock file is being used (see **use\_lockfile**), if a lock file already exists and is older than this value, it is assumed to have been left behind by accident, and Exim attempts to remove it.



<b>mailbox_filecount</b>	Use: <i>appendfile</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------------	------------------------	----------------------	-----------------------

If this option is set, it is expanded, and the result is taken as the current number of files in the mailbox. It must be a decimal number, optionally followed by K or M. This provides a way of obtaining this information from an external source that maintains the data.

<b>mailbox_size</b>	Use: <i>appendfile</i>	Type: <i>string†</i>	Default: <i>unset</i>
---------------------	------------------------	----------------------	-----------------------

If this option is set, it is expanded, and the result is taken as the current size the mailbox. It must be a decimal number, optionally followed by K or M. This provides a way of obtaining this information from an external source that maintains the data. This is likely to be helpful for maildir deliveries where it is computationally expensive to compute the size of a mailbox.

<b>maildir_format</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	------------------------	----------------------	-----------------------

If this option is set with the **directory** option, the delivery is into a new file, in the “maildir” format that is used by other mail software. When the transport is activated directly from a *redirect* router (for example, the *address\_file* transport in the default configuration), setting **maildir\_format** causes the path received from the router to be treated as a directory, whether or not it ends with /. This option is available only if SUPPORT\_MAILDIR is present in *Local/Makefile*. See section 26.5 below for further details.

<b>maildir_quota_directory_regex</b>	Use: <i>appendfile</i>	Type: <i>string</i>	Default: <i>See below</i>
--------------------------------------	------------------------	---------------------	---------------------------

This option is relevant only when **maildir\_use\_size\_file** is set. It defines a regular expression for specifying directories, relative to the quota directory (see **quota\_directory**), that should be included in the quota calculation. The default value is:

```
maildir_quota_directory_regex = ^(?:cur|new|\.*)$
```

This includes the *cur* and *new* directories, and any maildir++ folders (directories whose names begin with a dot). If you want to exclude the *Trash* folder from the count (as some sites do), you need to change this setting to

```
maildir_quota_directory_regex = ^(?:cur|new|\.(!Trash).*)$
```

This uses a negative lookahead in the regular expression to exclude the directory whose name is *.Trash*. When a directory is excluded from quota calculations, quota processing is bypassed for any messages that are delivered directly into that directory.

<b>maildir_retries</b>	Use: <i>appendfile</i>	Type: <i>integer</i>	Default: <i>10</i>
------------------------	------------------------	----------------------	--------------------

This option specifies the number of times to retry when writing a file in “maildir” format. See section 26.5 below.

<b>maildir_tag</b>	Use: <i>appendfile</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------	------------------------	----------------------	-----------------------

This option applies only to deliveries in maildir format, and is described in section 26.5 below.

<b>maildir_use_size_file</b>	Use: <i>appendfile†</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------------	-------------------------	----------------------	-----------------------

The result of string expansion for this option must be a valid boolean value. If it is true, it enables support for *maildirsize* files. Exim creates a *maildirsize* file in a maildir if one does not exist, taking the quota from the **quota** option of the transport. If **quota** is unset, the value is zero. See **maildir\_quota\_directory\_regex** above and section 26.5 below for further details.

<b>maildirfolder_create_regex</b>	Use: <i>appendfile</i>	Type: <i>string</i>	Default: <i>unset</i>
-----------------------------------	------------------------	---------------------	-----------------------

The value of this option is a regular expression. If it is *unset*, it has no effect. Otherwise, before a maildir delivery takes place, the pattern is matched against the name of the maildir directory, that is, the directory containing the *new* and *tmp* subdirectories that will be used for the delivery. If there is a match, Exim checks for the existence of a file called *maildirfolder* in the directory, and creates it if it does not exist. See section 26.5 for more details.

<b>mailstore_format</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------	------------------------	----------------------	-----------------------

If this option is set with the **directory** option, the delivery is into two new files in “mailstore” format. The option is available only if `SUPPORT_MAILSTORE` is present in *Local/Makefile*. See section 26.4 below for further details.

<b>mailstore_prefix</b>	Use: <i>appendfile</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-------------------------	------------------------	----------------------------------	-----------------------

This option applies only to deliveries in mailstore format, and is described in section 26.4 below.

<b>mailstore_suffix</b>	Use: <i>appendfile</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-------------------------	------------------------	----------------------------------	-----------------------

This option applies only to deliveries in mailstore format, and is described in section 26.4 below.

<b>mbx_format</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------	------------------------	----------------------	-----------------------

This option is available only if Exim has been compiled with `SUPPORT_MBX` set in *Local/Makefile*. If **mbx\_format** is set with the **file** option, the message is appended to the mailbox file in MBX format instead of traditional Unix format. This format is supported by Pine4 and its associated IMAP and POP daemons, by means of the *c-client* library that they all use.

**Note:** The **message\_prefix** and **message\_suffix** options are not automatically changed by the use of **mbx\_format**. They should normally be set empty when using MBX format, so this option almost always appears in this combination:

```
mbx_format = true
message_prefix =
message_suffix =
```

If none of the locking options are mentioned in the configuration, **use\_mbx\_lock** is assumed and the other locking options default to false. It is possible to specify the other kinds of locking with **mbx\_format**, but **use\_fcntl\_lock** and **use\_mbx\_lock** are mutually exclusive. MBX locking interworks with *c-client*, providing for shared access to the mailbox. It should not be used if any program that does not use this form of locking is going to access the mailbox, nor should it be used if the mailbox file is NFS mounted, because it works only when the mailbox is accessed from a single host.

If you set **use\_fcntl\_lock** with an MBX-format mailbox, you cannot use the standard version of *c-client*, because as long as it has a mailbox open (this means for the whole of a Pine or IMAP session), Exim will not be able to append messages to it.

<b>message_prefix</b>	Use: <i>appendfile</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
-----------------------	------------------------	----------------------------------	---------------------------

The string specified here is expanded and output at the start of every message. The default is *unset* unless **file** is specified and **use\_bsmtpl** is not set, in which case it is:

```
message_prefix = "From ${if def:return_path{$return_path}\
{MAILER-DAEMON}} $tod_bsdinbox\n"
```

**Note:** If you set **use\_crlf** true, you must change any occurrences of `\n` to `\r\n` in **message\_prefix**.

<b>message_suffix</b>	Use: <i>appendfile</i>	Type: <i>string†</i>	Default: <i>see below</i>
-----------------------	------------------------	----------------------	---------------------------

The string specified here is expanded and output at the end of every message. The default is unset unless **file** is specified and **use\_bsmtip** is not set, in which case it is a single newline character. The suffix can be suppressed by setting

```
message_suffix =
```

**Note:** If you set **use\_crlf** true, you must change any occurrences of `\n` to `\r\n` in **message\_suffix**.

<b>mode</b>	Use: <i>appendfile</i>	Type: <i>octal integer</i>	Default: <i>0600</i>
-------------	------------------------	----------------------------	----------------------

If the output file is created, it is given this mode. If it already exists and has wider permissions, they are reduced to this mode. If it has narrower permissions, an error occurs unless **mode\_fail\_narrower** is false. However, if the delivery is the result of a **save** command in a filter file specifying a particular mode, the mode of the output file is always forced to take that value, and this option is ignored.

<b>mode_fail_narrower</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------------------	------------------------	----------------------	----------------------

This option applies in the case when an existing mailbox file has a narrower mode than that specified by the **mode** option. If **mode\_fail\_narrower** is true, the delivery is deferred (“mailbox has the wrong mode”); otherwise Exim continues with the delivery attempt, using the existing mode of the file.

<b>notify_comsat</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------------	----------------------	-----------------------

If this option is true, the *comsat* daemon is notified after every successful delivery to a user mailbox. This is the daemon that notifies logged on users about incoming mail.

<b>quota</b>	Use: <i>appendfile</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------	------------------------	----------------------	-----------------------

This option imposes a limit on the size of the file to which Exim is appending, or to the total space used in the directory tree when the **directory** option is set. In the latter case, computation of the space used is expensive, because all the files in the directory (and any sub-directories) have to be individually inspected and their sizes summed. (See **quota\_size\_regex** and **maildir\_use\_size\_file** for ways to avoid this in environments where users have no shell access to their mailboxes).

As there is no interlock against two simultaneous deliveries into a multi-file mailbox, it is possible for the quota to be overrun in this case. For single-file mailboxes, of course, an interlock is a necessity.

A file’s size is taken as its *used* value. Because of blocking effects, this may be a lot less than the actual amount of disk space allocated to the file. If the sizes of a number of files are being added up, the rounding effect can become quite noticeable, especially on systems that have large block sizes. Nevertheless, it seems best to stick to the *used* figure, because this is the obvious value which users understand most easily.

The value of the option is expanded, and must then be a numerical value (decimal point allowed), optionally followed by one of the letters K, M, or G, for kilobytes, megabytes, or gigabytes, optionally followed by a slash and further option modifiers. If Exim is running on a system with large file support (Linux and FreeBSD have this), mailboxes larger than 2G can be handled.

The option modifier **no\_check** can be used to force delivery even if the over quota condition is met. The quota gets updated as usual.

**Note:** A value of zero is interpreted as “no quota”.

The expansion happens while Exim is running as root, before it changes uid for the delivery. This means that files that are inaccessible to the end user can be used to hold quota values that are looked up in the expansion. When delivery fails because this quota is exceeded, the handling of the error is as for system quota failures.

By default, Exim's quota checking mimics system quotas, and restricts the mailbox to the specified maximum size, though the value is not accurate to the last byte, owing to separator lines and additional headers that may get added during message delivery. When a mailbox is nearly full, large messages may get refused even though small ones are accepted, because the size of the current message is added to the quota when the check is made. This behaviour can be changed by setting **quota\_is\_inclusive** false. When this is done, the check for exceeding the quota does not include the current message. Thus, deliveries continue until the quota has been exceeded; thereafter, no further messages are delivered. See also **quota\_warn\_threshold**.

<b>quota_directory</b>	Use: <i>appendfile</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
------------------------	------------------------	----------------------------------	-----------------------

This option defines the directory to check for quota purposes when delivering into individual files. The default is the delivery directory, or, if a file called *maildirfolder* exists in a maildir directory, the parent of the delivery directory.

<b>quota_filecount</b>	Use: <i>appendfile</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>0</i>
------------------------	------------------------	----------------------------------	-------------------

This option applies when the **directory** option is set. It limits the total number of files in the directory (compare the inode limit in system quotas). It can only be used if **quota** is also set. The value is expanded; an expansion failure causes delivery to be deferred. A value of zero is interpreted as “no quota”.

The option modifier **no\_check** can be used to force delivery even if the over quota condition is met. The quota gets updated as usual.

<b>quota_is_inclusive</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------------------	------------------------	----------------------	----------------------

See **quota** above.

<b>quota_size_regex</b>	Use: <i>appendfile</i>	Type: <i>string</i>	Default: <i>unset</i>
-------------------------	------------------------	---------------------	-----------------------

This option applies when one of the delivery modes that writes a separate file for each message is being used. When Exim wants to find the size of one of these files in order to test the quota, it first checks **quota\_size\_regex**. If this is set to a regular expression that matches the file name, and it captures one string, that string is interpreted as a representation of the file's size. The value of **quota\_size\_regex** is not expanded.

This feature is useful only when users have no shell access to their mailboxes – otherwise they could defeat the quota simply by renaming the files. This facility can be used with maildir deliveries, by setting **maildir\_tag** to add the file length to the file name. For example:

```
maildir_tag = ,S=$message_size
quota_size_regex = ,S=(\d+)
```

An alternative to *\$message\_size* is *\$message\_linecount*, which contains the number of lines in the message.

The regular expression should not assume that the length is at the end of the file name (even though **maildir\_tag** puts it there) because maildir MUAs sometimes add other information onto the ends of message file names.

Section 26.7 contains further information.

<b>quota_warn_message</b>	Use: <i>appendfile</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
---------------------------	------------------------	----------------------------------	---------------------------

See below for the use of this option. If it is not set when **quota\_warn\_threshold** is set, it defaults to

```
quota_warn_message = "\
To: $local_part@$domain\n\
Subject: Your mailbox\n\n\
This message is automatically created \
by mail delivery software.\n\n\
The size of your mailbox has exceeded \
a warning threshold that is\n\
set by the system administrator.\n"
```

<b>quota_warn_threshold</b>	Use: <i>appendfile</i>	Type: <i>string†</i>	Default: <i>0</i>
-----------------------------	------------------------	----------------------	-------------------

This option is expanded in the same way as **quota** (see above). If the resulting value is greater than zero, and delivery of the message causes the size of the file or total space in the directory tree to cross the given threshold, a warning message is sent. If **quota** is also set, the threshold may be specified as a percentage of it by following the value with a percent sign. For example:

```
quota = 10M
quota_warn_threshold = 75%
```

If **quota** is not set, a setting of **quota\_warn\_threshold** that ends with a percent sign is ignored.

The warning message itself is specified by the **quota\_warn\_message** option, and it must start with a *To:* header line containing the recipient(s) of the warning message. These do not necessarily have to include the recipient(s) of the original message. A *Subject:* line should also normally be supplied. You can include any other header lines that you want. If you do not include a *From:* line, the default is:

```
From: Mail Delivery System <mailer-daemon@$qualify_domain_sender>
```

If you supply a *Reply-To:* line, it overrides the global **errors\_reply\_to** option.

The **quota** option does not have to be set in order to use this option; they are independent of one another except when the threshold is specified as a percentage.

<b>use_bsmtip</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------	------------------------	----------------------	-----------------------

If this option is set true, *appendfile* writes messages in “batch SMTP” format, with the envelope sender and recipient(s) included as SMTP commands. If you want to include a leading HELO command with such messages, you can do so by setting the **message\_prefix** option. See section 48.10 for details of batch SMTP.

<b>use_crlf</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------	------------------------	----------------------	-----------------------

This option causes lines to be terminated with the two-character CRLF sequence (carriage return, linefeed) instead of just a linefeed character. In the case of batched SMTP, the byte sequence written to the file is then an exact image of what would be sent down a real SMTP connection.

**Note:** The contents of the **message\_prefix** and **message\_suffix** options (which are used to supply the traditional “From ” and blank line separators in Berkeley-style mailboxes) are written verbatim, so must contain their own carriage return characters if these are needed. In cases where these options have non-empty defaults, the values end with a single linefeed, so they must be changed to end with `\r\n` if **use\_crlf** is set.

<b>use_fcntl_lock</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>see below</i>
-----------------------	------------------------	----------------------	---------------------------

This option controls the use of the *fcntl()* function to lock a file for exclusive use when a message is being appended. It is set by default unless **use\_flock\_lock** is set. Otherwise, it should be turned off only if you know that all your MUAs use lock file locking. When both **use\_fcntl\_lock** and **use\_flock\_lock** are unset, **use\_lockfile** must be set.

<b>use_flock_lock</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	------------------------	----------------------	-----------------------

This option is provided to support the use of *flock()* for file locking, for the few situations where it is needed. Most modern operating systems support *fcntl()* and *lockf()* locking, and these two functions interwork with each other. Exim uses *fcntl()* locking by default.

This option is required only if you are using an operating system where *flock()* is used by programs that access mailboxes (typically MUAs), and where *flock()* does not correctly interwork with *fcntl()*. You can use both *fcntl()* and *flock()* locking simultaneously if you want.

Not all operating systems provide *flock()*. Some versions of Solaris do not have it (and some, I think, provide a not quite right version built on top of *lockf()*). If the OS does not have *flock()*, Exim will be built without the ability to use it, and any attempt to do so will cause a configuration error.

**Warning:** *flock()* locks do not work on NFS files (unless *flock()* is just being mapped onto *fcntl()* by the OS).

<b>use_lockfile</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>see below</i>
---------------------	------------------------	----------------------	---------------------------

If this option is turned off, Exim does not attempt to create a lock file when appending to a mailbox file. In this situation, the only locking is by *fcntl()*. You should only turn **use\_lockfile** off if you are absolutely sure that every MUA that is ever going to look at your users' mailboxes uses *fcntl()* rather than a lock file, and even then only when you are not delivering over NFS from more than one host.

In order to append to an NFS file safely from more than one host, it is necessary to take out a lock *before* opening the file, and the lock file achieves this. Otherwise, even with *fcntl()* locking, there is a risk of file corruption.

The **use\_lockfile** option is set by default unless **use\_mbx\_lock** is set. It is not possible to turn both **use\_lockfile** and **use\_fcntl\_lock** off, except when **mbx\_format** is set.

<b>use_mbx_lock</b>	Use: <i>appendfile</i>	Type: <i>boolean</i>	Default: <i>see below</i>
---------------------	------------------------	----------------------	---------------------------

This option is available only if Exim has been compiled with **SUPPORT\_MBX** set in *Local/Makefile*. Setting the option specifies that special MBX locking rules be used. It is set by default if **mbx\_format** is set and none of the locking options are mentioned in the configuration. The locking rules are the same as are used by the *c-client* library that underlies Pine and the IMAP4 and POP daemons that come with it (see the discussion below). The rules allow for shared access to the mailbox. However, this kind of locking does not work when the mailbox is NFS mounted.

You can set **use\_mbx\_lock** with either (or both) of **use\_fcntl\_lock** and **use\_flock\_lock** to control what kind of locking is used in implementing the MBX locking rules. The default is to use *fcntl()* if **use\_mbx\_lock** is set without **use\_fcntl\_lock** or **use\_flock\_lock**.

## 26.3 Operational details for appending

Before appending to a file, the following preparations are made:

- If the name of the file is */dev/null*, no action is taken, and a success return is given.
- If any directories on the file's path are missing, Exim creates them if the **create\_directory** option is set. A created directory's mode is given by the **directory\_mode** option.
- If **file\_format** is set, the format of an existing file is checked. If this indicates that a different transport should be used, control is passed to that transport.
- If **use\_lockfile** is set, a lock file is built in a way that will work reliably over NFS, as follows:
  - (1) Create a "hitching post" file whose name is that of the lock file with the current time, primary host name, and process id added, by opening for writing as a new file. If this fails with an access error, delivery is deferred.

- (2) Close the hitching post file, and hard link it to the lock file name.
  - (3) If the call to *link()* succeeds, creation of the lock file has succeeded. Unlink the hitching post name.
  - (4) Otherwise, use *stat()* to get information about the hitching post file, and then unlink hitching post name. If the number of links is exactly two, creation of the lock file succeeded but something (for example, an NFS server crash and restart) caused this fact not to be communicated to the *link()* call.
  - (5) If creation of the lock file failed, wait for **lock\_interval** and try again, up to **lock\_retries** times. However, since any program that writes to a mailbox should complete its task very quickly, it is reasonable to time out old lock files that are normally the result of user agent and system crashes. If an existing lock file is older than **lockfile\_timeout** Exim attempts to unlink it before trying again.
- A call is made to *lstat()* to discover whether the main file exists, and if so, what its characteristics are. If *lstat()* fails for any reason other than non-existence, delivery is deferred.
  - If the file does exist and is a symbolic link, delivery is deferred, unless the **allow\_symlink** option is set, in which case the ownership of the link is checked, and then *stat()* is called to find out about the real file, which is then subjected to the checks below. The check on the top-level link ownership prevents one user creating a link for another's mailbox in a sticky directory, though allowing symbolic links in this case is definitely not a good idea. If there is a chain of symbolic links, the intermediate ones are not checked.
  - If the file already exists but is not a regular file, or if the file's owner and group (if the group is being checked – see **check\_group** above) are different from the user and group under which the delivery is running, delivery is deferred.
  - If the file's permissions are more generous than specified, they are reduced. If they are insufficient, delivery is deferred, unless **mode\_fail\_narrower** is set false, in which case the delivery is tried using the existing permissions.
  - The file's inode number is saved, and the file is then opened for appending. If this fails because the file has vanished, *appendfile* behaves as if it hadn't existed (see below). For any other failures, delivery is deferred.
  - If the file is opened successfully, check that the inode number hasn't changed, that it is still a regular file, and that the owner and permissions have not changed. If anything is wrong, defer delivery and freeze the message.
  - If the file did not exist originally, defer delivery if the **file\_must\_exist** option is set. Otherwise, check that the file is being created in a permitted directory if the **create\_file** option is set (deferring on failure), and then open for writing as a new file, with the O\_EXCL and O\_CREAT options, except when dealing with a symbolic link (the **allow\_symlink** option must be set). In this case, which can happen if the link points to a non-existent file, the file is opened for writing using O\_CREAT but not O\_EXCL, because that prevents link following.
  - If opening fails because the file exists, obey the tests given above for existing files. However, to avoid looping in a situation where the file is being continuously created and destroyed, the exists/not-exists loop is broken after 10 repetitions, and the message is then frozen.
  - If opening fails with any other error, defer delivery.
  - Once the file is open, unless both **use\_fcntl\_lock** and **use\_flock\_lock** are false, it is locked using *fcntl()* or *flock()* or both. If **use\_mbx\_lock** is false, an exclusive lock is requested in each case. However, if **use\_mbx\_lock** is true, Exim takes out a shared lock on the open file, and an exclusive lock on the file whose name is

/tmp/.<device-number>.<inode-number>

using the device and inode numbers of the open mailbox file, in accordance with the MBX locking rules. This file is created with a mode that is specified by the **lockfile\_mode** option.

If Exim fails to lock the file, there are two possible courses of action, depending on the value of the locking timeout. This is obtained from **lock\_fcntl\_timeout** or **lock\_flock\_timeout**, as appropriate.

If the timeout value is zero, the file is closed, Exim waits for **lock\_interval**, and then goes back and re-opens the file as above and tries to lock it again. This happens up to **lock\_retries** times, after which the delivery is deferred.

If the timeout has a value greater than zero, blocking calls to *fcntl()* or *flock()* are used (with the given timeout), so there has already been some waiting involved by the time locking fails. Nevertheless, Exim does not give up immediately. It retries up to

$$(\text{lock\_retries} * \text{lock\_interval}) / <\text{timeout}>$$

times (rounded up).

At the end of delivery, Exim closes the file (which releases the *fcntl()* and/or *flock()* locks) and then deletes the lock file if one was created.

## 26.4 Operational details for delivery to a new file

When the **directory** option is set instead of **file**, each message is delivered into a newly-created file or set of files. When *appendfile* is activated directly from a *redirect* router, neither **file** nor **directory** is normally set, because the path for delivery is supplied by the router. (See for example, the *address\_file* transport in the default configuration.) In this case, delivery is to a new file if either the path name ends in /, or the **maildir\_format** or **mailstore\_format** option is set.

No locking is required while writing the message to a new file, so the various locking options of the transport are ignored. The “From” line that by default separates messages in a single file is not normally needed, nor is the escaping of message lines that start with “From”, and there is no need to ensure a newline at the end of each message. Consequently, the default values for **check\_string**, **message\_prefix**, and **message\_suffix** are all unset when any of **directory**, **maildir\_format**, or **mailstore\_format** is set.

If Exim is required to check a **quota** setting, it adds up the sizes of all the files in the delivery directory by default. However, you can specify a different directory by setting **quota\_directory**. Also, for maildir deliveries (see below) the *maildirfolder* convention is honoured.

There are three different ways in which delivery to individual files can be done, controlled by the settings of the **maildir\_format** and **mailstore\_format** options. Note that code to support maildir or mailstore formats is not included in the binary unless **SUPPORT\_MAILDIR** or **SUPPORT\_MAILSTORE**, respectively, is set in *Local/Makefile*.

In all three cases an attempt is made to create the directory and any necessary sub-directories if they do not exist, provided that the **create\_directory** option is set (the default). The location of a created directory can be constrained by setting **create\_file**. A created directory’s mode is given by the **directory\_mode** option. If creation fails, or if the **create\_directory** option is not set when creation is required, delivery is deferred.

## 26.5 Maildir delivery

If the **maildir\_format** option is true, Exim delivers each message by writing it to a file whose name is *tmp/<stime>.H<mtime>P<pid>.<host>* in the directory that is defined by the **directory** option (the “delivery directory”). If the delivery is successful, the file is renamed into the *new* subdirectory.

In the file name, *<stime>* is the current time of day in seconds, and *<mtime>* is the microsecond fraction of the time. After a maildir delivery, Exim checks that the time-of-day clock has moved on by at least one microsecond before terminating the delivery process. This guarantees uniqueness for the file name. However, as a precaution, Exim calls *stat()* for the file before opening it. If any response other than ENOENT (does not exist) is given, Exim waits 2 seconds and tries again, up to **maildir\_retries** times.

Before Exim carries out a maildir delivery, it ensures that subdirectories called *new*, *cur*, and *tmp* exist in the delivery directory. If they do not exist, Exim tries to create them and any superior directories in



their path, subject to the **create\_directory** and **create\_file** options. If the **maildirfolder\_create\_regex** option is set, and the regular expression it contains matches the delivery directory, Exim also ensures that a file called *maildirfolder* exists in the delivery directory. If a missing directory or *maildirfolder* file cannot be created, delivery is deferred.

These features make it possible to use Exim to create all the necessary files and directories in a maildir mailbox, including subdirectories for maildir++ folders. Consider this example:

```
maildir_format = true
directory = /var/mail/$local_part\
           ${if eq{$local_part_suffix}{}{}}\
           {/.${substr_1:$local_part_suffix}}
maildirfolder_create_regex = /\.[^/]+$
```

If *\$local\_part\_suffix* is empty (there was no suffix for the local part), delivery is into a toplevel maildir with a name like */var/mail/pimbo* (for the user called *pimbo*). The pattern in **maildirfolder\_create\_regex** does not match this name, so Exim will not look for or create the file */var/mail/pimbo/maildirfolder*, though it will create */var/mail/pimbo/{cur,new,tmp}* if necessary.

However, if *\$local\_part\_suffix* contains *-eximusers* (for example), delivery is into the maildir++ folder */var/mail/pimbo.eximusers*, which does match **maildirfolder\_create\_regex**. In this case, Exim will create */var/mail/pimbo.eximusers/maildirfolder* as well as the three maildir directories */var/mail/pimbo.eximusers/{cur,new,tmp}*.

**Warning:** Take care when setting **maildirfolder\_create\_regex** that it does not inadvertently match the toplevel maildir directory, because a *maildirfolder* file at top level would completely break quota calculations.

If Exim is required to check a **quota** setting before a maildir delivery, and **quota\_directory** is not set, it looks for a file called *maildirfolder* in the maildir directory (alongside *new*, *cur*, *tmp*). If this exists, Exim assumes the directory is a maildir++ folder directory, which is one level down from the user's top level mailbox directory. This causes it to start at the parent directory instead of the current directory when calculating the amount of space used.

One problem with delivering into a multi-file mailbox is that it is computationally expensive to compute the size of the mailbox for quota checking. Various approaches have been taken to reduce the amount of work needed. The next two sections describe two of them. A third alternative is to use some external process for maintaining the size data, and use the expansion of the **mailbox\_size** option as a way of importing it into Exim.

## 26.6 Using tags to record message sizes

If **maildir\_tag** is set, the string is expanded for each delivery. When the maildir file is renamed into the *new* sub-directory, the tag is added to its name. However, if adding the tag takes the length of the name to the point where the test *stat()* call fails with ENAMETOOLONG, the tag is dropped and the maildir file is created with no tag.

Tags can be used to encode the size of files in their names; see **quota\_size\_regex** above for an example. The expansion of **maildir\_tag** happens after the message has been written. The value of the *\$message\_size* variable is set to the number of bytes actually written. If the expansion is forced to fail, the tag is ignored, but a non-forced failure causes delivery to be deferred. The expanded tag may contain any printing characters except */*. Non-printing characters in the string are ignored; if the resulting string is empty, it is ignored. If it starts with an alphanumeric character, a leading colon is inserted; this default has not proven to be the path that popular maildir implementations have chosen (but changing it in Exim would break backwards compatibility).

For one common implementation, you might set:

```
maildir_tag = ,S=${message_size}
```

but you should check the documentation of the other software to be sure.

It is advisable to also set **quota\_size\_regex** when setting **maildir\_tag** as this allows Exim to extract the size from your tag, instead of having to *stat()* each message file.

## 26.7 Using a maildirsize file

If **maildir\_use\_size\_file** is true, Exim implements the maildir++ rules for storing quota and message size information in a file called *maildirsize* within the toplevel maildir directory. If this file does not exist, Exim creates it, setting the quota from the **quota** option of the transport. If the maildir directory itself does not exist, it is created before any attempt to write a *maildirsize* file.

The *maildirsize* file is used to hold information about the sizes of messages in the maildir, thus speeding up quota calculations. The quota value in the file is just a cache; if the quota is changed in the transport, the new value overrides the cached value when the next message is delivered. The cache is maintained for the benefit of other programs that access the maildir and need to know the quota.

If the **quota** option in the transport is unset or zero, the *maildirsize* file is maintained (with a zero quota setting), but no quota is imposed.

A regular expression is available for controlling which directories in the maildir participate in quota calculations when a *maildirsizefile* is in use. See the description of the **maildir\_quota\_directory\_regex** option above for details.

## 26.8 Mailstore delivery

If the **mailstore\_format** option is true, each message is written as two files in the given directory. A unique base name is constructed from the message id and the current delivery process, and the files that are written use this base name plus the suffixes *.env* and *.msg*. The *.env* file contains the message's envelope, and the *.msg* file contains the message itself. The base name is placed in the variable *\$mailstore\_basename*.

During delivery, the envelope is first written to a file with the suffix *.tmp*. The *.msg* file is then written, and when it is complete, the *.tmp* file is renamed as the *.env* file. Programs that access messages in mailstore format should wait for the presence of both a *.msg* and a *.env* file before accessing either of them. An alternative approach is to wait for the absence of a *.tmp* file.

The envelope file starts with any text defined by the **mailstore\_prefix** option, expanded and terminated by a newline if there isn't one. Then follows the sender address on one line, then all the recipient addresses, one per line. There can be more than one recipient only if the **batch\_max** option is set greater than one. Finally, **mailstore\_suffix** is expanded and the result appended to the file, followed by a newline if it does not end with one.

If expansion of **mailstore\_prefix** or **mailstore\_suffix** ends with a forced failure, it is ignored. Other expansion errors are treated as serious configuration errors, and delivery is deferred. The variable *\$mailstore\_basename* is available for use during these expansions.

## 26.9 Non-special new file delivery

If neither **maildir\_format** nor **mailstore\_format** is set, a single new file is created directly in the named directory. For example, when delivering messages into files in batched SMTP format for later delivery to some host (see section 48.10), a setting such as

```
directory = /var/bsmtp/$host
```

might be used. A message is written to a file with a temporary name, which is then renamed when the delivery is complete. The final name is obtained by expanding the contents of the **directory\_file** option.

## 27. The autoreply transport

The *autoreply* transport is not a true transport in that it does not cause the message to be transmitted. Instead, it generates a new mail message as an automatic reply to the incoming message. *References:* and *Auto-Submitted:* header lines are included. These are constructed according to the rules in RFCs 2822 and 3834, respectively.

If the router that passes the message to this transport does not have the **unseen** option set, the original message (for the current recipient) is not delivered anywhere. However, when the **unseen** option is set on the router that passes the message to this transport, routing of the address continues, so another router can set up a normal message delivery.

The *autoreply* transport is usually run as the result of mail filtering, a “vacation” message being the standard example. However, it can also be run directly from a router like any other transport. To reduce the possibility of message cascades, messages created by the *autoreply* transport always have empty envelope sender addresses, like bounce messages.

The parameters of the message to be sent can be specified in the configuration by options described below. However, these are used only when the address passed to the transport does not contain its own reply information. When the transport is run as a consequence of a **mail** or **vacation** command in a filter file, the parameters of the message are supplied by the filter, and passed with the address. The transport’s options that define the message are then ignored (so they are not usually set in this case). The message is specified entirely by the filter or by the transport; it is never built from a mixture of options. However, the **file\_optional**, **mode**, and **return\_message** options apply in all cases.

*Autoreply* is implemented as a local transport. When used as a result of a command in a user’s filter file, *autoreply* normally runs under the uid and gid of the user, and with appropriate current and home directories (see chapter 23).

There is a subtle difference between routing a message to a *pipe* transport that generates some text to be returned to the sender, and routing it to an *autoreply* transport. This difference is noticeable only if more than one address from the same message is so handled. In the case of a pipe, the separate outputs from the different addresses are gathered up and returned to the sender in a single message, whereas if *autoreply* is used, a separate message is generated for each address that is passed to it.

Non-printing characters are not permitted in the header lines generated for the message that *autoreply* creates, with the exception of newlines that are immediately followed by white space. If any non-printing characters are found, the transport defers. Whether characters with the top bit set count as printing characters or not is controlled by the **print\_topbitchars** global option.

If any of the generic options for manipulating headers (for example, **headers\_add**) are set on an *autoreply* transport, they apply to the copy of the original message that is included in the generated message when **return\_message** is set. They do not apply to the generated message itself.

If the *autoreply* transport receives return code 2 from Exim when it submits the message, indicating that there were no recipients, it does not treat this as an error. This means that autoreplies sent to *\$sender\_address* when this is empty (because the incoming message is a bounce message) do not cause problems. They are just discarded.

### 27.1 Private options for autoreply

<b>bcc</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
------------	-----------------------	----------------------------------	-----------------------

This specifies the addresses that are to receive “blind carbon copies” of the message when the message is specified by the transport.

<b>cc</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-----------	-----------------------	----------------------------------	-----------------------

This specifies recipients of the message and the contents of the *Cc:* header when the message is specified by the transport.

<b>file</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-------------	-----------------------	----------------------------------	-----------------------

The contents of the file are sent as the body of the message when the message is specified by the transport. If both **file** and **text** are set, the text string comes first.

<b>file_expand</b>	Use: <i>autoreply</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------	-----------------------	----------------------	-----------------------

If this is set, the contents of the file named by the **file** option are subjected to string expansion as they are added to the message.

<b>file_optional</b>	Use: <i>autoreply</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	-----------------------	----------------------	-----------------------

If this option is true, no error is generated if the file named by the **file** option or passed with the address does not exist or cannot be read.

<b>from</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-------------	-----------------------	----------------------------------	-----------------------

This specifies the contents of the *From:* header when the message is specified by the transport.

<b>headers</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
----------------	-----------------------	----------------------------------	-----------------------

This specifies additional RFC 2822 headers that are to be added to the message when the message is specified by the transport. Several can be given by using “\n” to separate them. There is no check on the format.

<b>log</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
------------	-----------------------	----------------------------------	-----------------------

This option names a file in which a record of every message sent is logged when the message is specified by the transport.

<b>mode</b>	Use: <i>autoreply</i>	Type: <i>octal integer</i>	Default: <i>0600</i>
-------------	-----------------------	----------------------------	----------------------

If either the log file or the “once” file has to be created, this mode is used.

<b>never_mail</b>	Use: <i>autoreply</i>	Type: <i>address list</i> <sup>†</sup>	Default: <i>unset</i>
-------------------	-----------------------	--	-----------------------

If any run of the transport creates a message with a recipient that matches any item in the list, that recipient is quietly discarded. If all recipients are discarded, no message is created. This applies both when the recipients are generated by a filter and when they are specified in the transport.

<b>once</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-------------	-----------------------	----------------------------------	-----------------------

This option names a file or DBM database in which a record of each *To:* recipient is kept when the message is specified by the transport. **Note:** This does not apply to *Cc:* or *Bcc:* recipients.

If **once** is unset, or is set to an empty string, the message is always sent. By default, if **once** is set to a non-empty file name, the message is not sent if a potential recipient is already listed in the database. However, if the **once\_repeat** option specifies a time greater than zero, the message is sent if that

much time has elapsed since a message was last sent to this recipient. A setting of zero time for **once\_repeat** (the default) prevents a message from being sent a second time – in this case, zero means infinity.

If **once\_file\_size** is zero, a DBM database is used to remember recipients, and it is allowed to grow as large as necessary. If **once\_file\_size** is set greater than zero, it changes the way Exim implements the **once** option. Instead of using a DBM file to record every recipient it sends to, it uses a regular file, whose size will never get larger than the given value.

In the file, Exim keeps a linear list of recipient addresses and the times at which they were sent messages. If the file is full when a new address needs to be added, the oldest address is dropped. If **once\_repeat** is not set, this means that a given recipient may receive multiple messages, but at unpredictable intervals that depend on the rate of turnover of addresses in the file. If **once\_repeat** is set, it specifies a maximum time between repeats.

<b>once_file_size</b>	Use: <i>autoreply</i>	Type: <i>integer</i>	Default: <i>0</i>
-----------------------	-----------------------	----------------------	-------------------

See **once** above.

<b>once_repeat</b>	Use: <i>autoreply</i>	Type: <i>time</i> <sup>†</sup>	Default: <i>0s</i>
--------------------	-----------------------	--------------------------------	--------------------

See **once** above. After expansion, the value of this option must be a valid time value.

<b>reply_to</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-----------------	-----------------------	----------------------------------	-----------------------

This specifies the contents of the *Reply-To:* header when the message is specified by the transport.

<b>return_message</b>	Use: <i>autoreply</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	-----------------------	----------------------	-----------------------

If this is set, a copy of the original message is returned with the new message, subject to the maximum size set in the **return\_size\_limit** global configuration option.

<b>subject</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
----------------	-----------------------	----------------------------------	-----------------------

This specifies the contents of the *Subject:* header when the message is specified by the transport. It is tempting to quote the original subject in automatic responses. For example:

```
subject = Re: $h_subject:
```

There is a danger in doing this, however. It may allow a third party to subscribe your users to an opt-in mailing list, provided that the list accepts bounce messages as subscription confirmations. Well-managed lists require a non-bounce message to confirm a subscription, so the danger is relatively small.

<b>text</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-------------	-----------------------	----------------------------------	-----------------------

This specifies a single string to be used as the body of the message when the message is specified by the transport. If both **text** and **file** are set, the text comes first.

<b>to</b>	Use: <i>autoreply</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-----------	-----------------------	----------------------------------	-----------------------

This specifies recipients of the message and the contents of the *To:* header when the message is specified by the transport.

## 28. The lmtip transport

The *lmtip* transport runs the LMTP protocol (RFC 2033) over a pipe to a specified command or by interacting with a Unix domain socket. This transport is something of a cross between the *pipe* and *smtp* transports. Exim also has support for using LMTP over TCP/IP; this is implemented as an option for the *smtp* transport. Because LMTP is expected to be of minority interest, the default build-time configure in *src/EDITME* has it commented out. You need to ensure that

```
TRANSPORT_LMTP=yes
```

is present in your *Local/Makefile* in order to have the *lmtip* transport included in the Exim binary. The private options of the *lmtip* transport are as follows:

<b>batch_id</b>	Use: <i>lmtip</i>	Type: <i>string†</i>	Default: <i>unset</i>
-----------------	-------------------	----------------------	-----------------------

See the description of local delivery batching in chapter 25.

<b>batch_max</b>	Use: <i>lmtip</i>	Type: <i>integer</i>	Default: <i>1</i>
------------------	-------------------	----------------------	-------------------

This limits the number of addresses that can be handled in a single delivery. Most LMTP servers can handle several addresses at once, so it is normally a good idea to increase this value. See the description of local delivery batching in chapter 25.

<b>command</b>	Use: <i>lmtip</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------	-------------------	----------------------	-----------------------

This option must be set if **socket** is not set. The string is a command which is run in a separate process. It is split up into a command name and list of arguments, each of which is separately expanded (so expansion cannot change the number of arguments). The command is run directly, not via a shell. The message is passed to the new process using the standard input and output to operate the LMTP protocol.

<b>ignore_quota</b>	Use: <i>lmtip</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------	-------------------	----------------------	-----------------------

If this option is set true, the string `IGNOREQUOTA` is added to RCPT commands, provided that the LMTP server has advertised support for `IGNOREQUOTA` in its response to the LHLO command.

<b>socket</b>	Use: <i>lmtip</i>	Type: <i>string†</i>	Default: <i>unset</i>
---------------	-------------------	----------------------	-----------------------

This option must be set if **command** is not set. The result of expansion must be the name of a Unix domain socket. The transport connects to the socket and delivers the message to it using the LMTP protocol.

<b>timeout</b>	Use: <i>lmtip</i>	Type: <i>time</i>	Default: <i>5m</i>
----------------	-------------------	-------------------	--------------------

The transport is aborted if the created process or Unix domain socket does not respond to LMTP commands or message input within this timeout. Delivery is deferred, and will be tried again later. Here is an example of a typical LMTP transport:

```
lmtip:
  driver = lmtip
  command = /some/local/lmtip/delivery/program
  batch_max = 20
  user = exim
```

This delivers up to 20 addresses at a time, in a mixture of domains if necessary, running as the user *exim*.

## 29. The pipe transport

The *pipe* transport is used to deliver messages via a pipe to a command running in another process. One example is the use of *pipe* as a pseudo-remote transport for passing messages to some other delivery mechanism (such as UUCP). Another is the use by individual users to automatically process their incoming messages. The *pipe* transport can be used in one of the following ways:

- A router routes one address to a transport in the normal way, and the transport is configured as a *pipe* transport. In this case, *\$local\_part* contains the local part of the address (as usual), and the command that is run is specified by the **command** option on the transport.
- If the **batch\_max** option is set greater than 1 (the default is 1), the transport can handle more than one address in a single run. In this case, when more than one address is routed to the transport, *\$local\_part* is not set (because it is not unique). However, the pseudo-variable *\$pipe\_addresses* (described in section 29.3 below) contains all the addresses that are routed to the transport.
- A router redirects an address directly to a pipe command (for example, from an alias or forward file). In this case, *\$address\_pipe* contains the text of the pipe command, and the **command** option on the transport is ignored unless **force\_command** is set. If only one address is being transported (**batch\_max** is not greater than one, or only one address was redirected to this pipe command), *\$local\_part* contains the local part that was redirected.

The *pipe* transport is a non-interactive delivery method. Exim can also deliver messages over pipes using the LMTP interactive protocol. This is implemented by the *lmtp* transport.

In the case when *pipe* is run as a consequence of an entry in a local user's *forward* file, the command runs under the uid and gid of that user. In other cases, the uid and gid have to be specified explicitly, either on the transport or on the router that handles the address. Current and "home" directories are also controllable. See chapter 23 for details of the local delivery environment and chapter 25 for a discussion of local delivery batching.

### 29.1 Concurrent delivery

If two messages arrive at almost the same time, and both are routed to a pipe delivery, the two pipe transports may be run concurrently. You must ensure that any pipe commands you set up are robust against this happening. If the commands write to a file, the **exim\_lock** utility might be of use. Alternatively the **max\_parallel** option could be used with a value of "1" to enforce serialization.

### 29.2 Returned status and data

If the command exits with a non-zero return code, the delivery is deemed to have failed, unless either the **ignore\_status** option is set (in which case the return code is treated as zero), or the return code is one of those listed in the **temp\_errors** option, which are interpreted as meaning "try again later". In this case, delivery is deferred. Details of a permanent failure are logged, but are not included in the bounce message, which merely contains "local delivery failed".

If the command exits on a signal and the **freeze\_signal** option is set then the message will be frozen in the queue. If that option is not set, a bounce will be sent as normal.

If the return code is greater than 128 and the command being run is a shell script, it normally means that the script was terminated by a signal whose value is the return code minus 128. The **freeze\_signal** option does not apply in this case.

If Exim is unable to run the command (that is, if *execve()* fails), the return code is set to 127. This is the value that a shell returns if it is asked to run a non-existent command. The wording for the log line suggests that a non-existent command may be the problem.

The **return\_output** option can affect the result of a pipe delivery. If it is set and the command produces any output on its standard output or standard error streams, the command is considered to have failed, even if it gave a zero return code or if **ignore\_status** is set. The output from the command is included as part of the bounce message. The **return\_fail\_output** option is similar, except that



output is returned only when the command exits with a failure return code, that is, a value other than zero or a code that matches **temp\_errors**.

## 29.3 How the command is run

The command line is (by default) broken down into a command name and arguments by the *pipe* transport itself. The **allow\_commands** and **restrict\_to\_path** options can be used to restrict the commands that may be run.

Unquoted arguments are delimited by white space. If an argument appears in double quotes, backslash is interpreted as an escape character in the usual way. If an argument appears in single quotes, no escaping is done.

String expansion is applied to the command line except when it comes from a traditional *forward* file (commands from a filter file are expanded). The expansion is applied to each argument in turn rather than to the whole line. For this reason, any string expansion item that contains white space must be quoted so as to be contained within a single argument. A setting such as

```
command = /some/path ${if eq{$local_part}{postmaster}{xx}{yy}}
```

will not work, because the expansion item gets split between several arguments. You have to write

```
command = /some/path "${if eq{$local_part}{postmaster}{xx}{yy}}"
```

to ensure that it is all in one argument. The expansion is done in this way, argument by argument, so that the number of arguments cannot be changed as a result of expansion, and quotes or backslashes in inserted variables do not interact with external quoting. However, this leads to problems if you want to generate multiple arguments (or the command name plus arguments) from a single expansion. In this situation, the simplest solution is to use a shell. For example:

```
command = /bin/sh -c ${lookup{$local_part}lsearch{/some/file}}
```

Special handling takes place when an argument consists of precisely the text `$pipe_addresses`. This is not a general expansion variable; the only place this string is recognized is when it appears as an argument for a pipe or transport filter command. It causes each address that is being handled to be inserted in the argument list at that point *as a separate argument*. This avoids any problems with spaces or shell metacharacters, and is of use when a *pipe* transport is handling groups of addresses in a batch.

If **force\_command** is enabled on the transport, Special handling takes place for an argument that consists of precisely the text `$address_pipe`. It is handled similarly to `$pipe_addresses` above. It is expanded and each argument is inserted in the argument list at that point *as a separate argument*. The `$address_pipe` item does not need to be the only item in the argument; in fact, if it were then **force\_command** should behave as a no-op. Rather, it should be used to adjust the command run while preserving the argument vector separation.

After splitting up into arguments and expansion, the resulting command is run in a subprocess directly from the transport, *not* under a shell. The message that is being delivered is supplied on the standard input, and the standard output and standard error are both connected to a single pipe that is read by Exim. The **max\_output** option controls how much output the command may produce, and the **return\_output** and **return\_fail\_output** options control what is done with it.

Not running the command under a shell (by default) lessens the security risks in cases when a command from a user's filter file is built out of data that was taken from an incoming message. If a shell is required, it can of course be explicitly specified as the command to be run. However, there are circumstances where existing commands (for example, in *forward* files) expect to be run under a shell and cannot easily be modified. To allow for these cases, there is an option called **use\_shell**, which changes the way the *pipe* transport works. Instead of breaking up the command line as just described, it expands it as a single string and passes the result to `/bin/sh`. The **restrict\_to\_path** option and the `$pipe_addresses` facility cannot be used with **use\_shell**, and the whole mechanism is inherently less secure.

## 29.4 Environment variables

The environment variables listed below are set up when the command is invoked. This list is a compromise for maximum compatibility with other MTAs. Note that the **environment** option can be used to add additional variables to this environment. The environment for the *pipe* transport is not subject to the **add\_environment** and **keep\_environment** main config options.

DOMAIN	the domain of the address
HOME	the home directory, if set
HOST	the host name when called from a router (see below)
LOCAL_PART	see below
LOCAL_PART_PREFIX	see below
LOCAL_PART_SUFFIX	see below
LOGNAME	see below
MESSAGE_ID	Exim's local ID for the message
PATH	as specified by the <b>path</b> option below
QUALIFY_DOMAIN	the sender qualification domain
RECIPIENT	the complete recipient address
SENDER	the sender of the message (empty if a bounce)
SHELL	/bin/sh
TZ	the value of the <b>timezone</b> option, if set
USER	see below

When a *pipe* transport is called directly from (for example) an *accept* router, **LOCAL\_PART** is set to the local part of the address. When it is called as a result of a forward or alias expansion, **LOCAL\_PART** is set to the local part of the address that was expanded. In both cases, any affixes are removed from the local part, and made available in **LOCAL\_PART\_PREFIX** and **LOCAL\_PART\_SUFFIX**, respectively. **LOGNAME** and **USER** are set to the same value as **LOCAL\_PART** for compatibility with other MTAs.

**HOST** is set only when a *pipe* transport is called from a router that associates hosts with an address, typically when using *pipe* as a pseudo-remote transport. **HOST** is set to the first host name specified by the router.

If the transport's generic **home\_directory** option is set, its value is used for the **HOME** environment variable. Otherwise, a home directory may be set by the router's **transport\_home\_directory** option, which defaults to the user's home directory if **check\_local\_user** is set.

## 29.5 Private options for pipe

<b>allow_commands</b>	Use: <i>pipe</i>	Type: <i>string list</i> <sup>†</sup>	Default: <i>unset</i>
-----------------------	------------------	---------------------------------------	-----------------------

The string is expanded, and is then interpreted as a colon-separated list of permitted commands. If **restrict\_to\_path** is not set, the only commands permitted are those in the **allow\_commands** list. They need not be absolute paths; the **path** option is still used for relative paths. If **restrict\_to\_path** is set with **allow\_commands**, the command must either be in the **allow\_commands** list, or a name without any slashes that is found on the path. In other words, if neither **allow\_commands** nor **restrict\_to\_path** is set, there is no restriction on the command, but otherwise only commands that are permitted by one or the other are allowed. For example, if

```
allow_commands = /usr/bin/vacation
```

and **restrict\_to\_path** is not set, the only permitted command is */usr/bin/vacation*. The **allow\_commands** option may not be set if **use\_shell** is set.

<b>batch_id</b>	Use: <i>pipe</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-----------------	------------------	----------------------------------	-----------------------

See the description of local delivery batching in chapter 25.

<b>batch_max</b>	Use: <i>pipe</i>	Type: <i>integer</i>	Default: <i>1</i>
------------------	------------------	----------------------	-------------------

This limits the number of addresses that can be handled in a single delivery. See the description of local delivery batching in chapter 25.

<b>check_string</b>	Use: <i>pipe</i>	Type: <i>string</i>	Default: <i>unset</i>
---------------------	------------------	---------------------	-----------------------

As *pipe* writes the message, the start of each line is tested for matching **check\_string**, and if it does, the initial matching characters are replaced by the contents of **escape\_string**, provided both are set. The value of **check\_string** is a literal string, not a regular expression, and the case of any letters it contains is significant. When **use\_bsmtplib** is set, the contents of **check\_string** and **escape\_string** are forced to values that implement the SMTP escaping protocol. Any settings made in the configuration file are ignored.

<b>command</b>	Use: <i>pipe</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
----------------	------------------	----------------------------------	-----------------------

This option need not be set when *pipe* is being used to deliver to pipes obtained directly from address redirections. In other cases, the option must be set, to provide a command to be run. It need not yield an absolute path (see the **path** option below). The command is split up into separate arguments by Exim, and each argument is separately expanded, as described in section 29.3 above.

<b>environment</b>	Use: <i>pipe</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
--------------------	------------------	----------------------------------	-----------------------

This option is used to add additional variables to the environment in which the command runs (see section 29.4 for the default list). Its value is a string which is expanded, and then interpreted as a colon-separated list of environment settings of the form `<name>=<value>`.

<b>escape_string</b>	Use: <i>pipe</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------	------------------	---------------------	-----------------------

See **check\_string** above.

<b>freeze_exec_fail</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------	------------------	----------------------	-----------------------

Failure to exec the command in a pipe transport is by default treated like any other failure while running the command. However, if **freeze\_exec\_fail** is set, failure to exec is treated specially, and causes the message to be frozen, whatever the setting of **ignore\_status**.

<b>freeze_signal</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------	----------------------	-----------------------

Normally if the process run by a command in a pipe transport exits on a signal, a bounce message is sent. If **freeze\_signal** is set, the message will be frozen in Exim's queue instead.

<b>force_command</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------	----------------------	-----------------------

Normally when a router redirects an address directly to a pipe command the **command** option on the transport is ignored. If **force\_command** is set, the **command** option will be used. This is especially useful for forcing a wrapper or additional argument to be added to the command. For example:

```
command = /usr/bin/remote_exec myhost -- $address_pipe
force_command
```

Note that *\$address\_pipe* is handled specially in **command** when **force\_command** is set, expanding out to the original argument vector as separate items, similarly to a Unix shell "`$@`" construct.

<b>ignore_status</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------	----------------------	-----------------------

If this option is true, the status returned by the subprocess that is set up to run the command is ignored, and Exim behaves as if zero had been returned. Otherwise, a non-zero status or termination by signal causes an error return from the transport unless the status value is one of those listed in **temp\_errors**; these cause the delivery to be deferred and tried again later.

**Note:** This option does not apply to timeouts, which do not return a status. See the **timeout\_defer** option for how timeouts are handled.

<b>log_defer_output</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------	------------------	----------------------	-----------------------

If this option is set, and the status returned by the command is one of the codes listed in **temp\_errors** (that is, delivery was deferred), and any output was produced on stdout or stderr, the first line of it is written to the main log.

<b>log_fail_output</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------	----------------------	-----------------------

If this option is set, and the command returns any output on stdout or stderr, and also ends with a return code that is neither zero nor one of the return codes listed in **temp\_errors** (that is, the delivery failed), the first line of output is written to the main log. This option and **log\_output** are mutually exclusive. Only one of them may be set.

<b>log_output</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------	------------------	----------------------	-----------------------

If this option is set and the command returns any output on stdout or stderr, the first line of output is written to the main log, whatever the return code. This option and **log\_fail\_output** are mutually exclusive. Only one of them may be set.

<b>max_output</b>	Use: <i>pipe</i>	Type: <i>integer</i>	Default: <i>20K</i>
-------------------	------------------	----------------------	---------------------

This specifies the maximum amount of output that the command may produce on its standard output and standard error file combined. If the limit is exceeded, the process running the command is killed. This is intended as a safety measure to catch runaway processes. The limit is applied independently of the settings of the options that control what is done with such output (for example, **return\_output**). Because of buffering effects, the amount of output may exceed the limit by a small amount before Exim notices.

<b>message_prefix</b>	Use: <i>pipe</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
-----------------------	------------------	----------------------------------	---------------------------

The string specified here is expanded and output at the start of every message. The default is unset if **use\_bsmtplib** is set. Otherwise it is

```
message_prefix = \
  From ${if def:return_path{$return_path}{MAILER-DAEMON}} \
  ${tod_bsdinbox}\n
```

This is required by the commonly used */usr/bin/vacation* program. However, it must *not* be present if delivery is to the Cyrus IMAP server, or to the **tmllib** local delivery agent. The prefix can be suppressed by setting

```
message_prefix =
```

**Note:** If you set **use\_crlf** true, you must change any occurrences of `\n` to `\r\n` in **message\_prefix**.

<b>message_suffix</b>	Use: <i>pipe</i>	Type: <i>string†</i>	Default: <i>see below</i>
-----------------------	------------------	----------------------	---------------------------

The string specified here is expanded and output at the end of every message. The default is unset if **use\_bsmtplib** is set. Otherwise it is a single newline. The suffix can be suppressed by setting

```
message_suffix =
```

**Note:** If you set **use\_crlf** true, you must change any occurrences of `\n` to `\r\n` in **message\_suffix**.

<b>path</b>	Use: <i>pipe</i>	Type: <i>string†</i>	Default: <i>/bin:/usr/bin</i>
-------------	------------------	----------------------	-------------------------------

This option is expanded and specifies the string that is set up in the PATH environment variable of the subprocess. If the **command** option does not yield an absolute path name, the command is sought in the PATH directories, in the usual way. **Warning:** This does not apply to a command specified as a transport filter.

<b>permit_coredump</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------	----------------------	-----------------------

Normally Exim inhibits core-dumps during delivery. If you have a need to get a core-dump of a pipe command, enable this command. This enables core-dumps during delivery and affects both the Exim binary and the pipe command run. It is recommended that this option remain off unless and until you have a need for it and that this only be enabled when needed, as the risk of excessive resource consumption can be quite high. Note also that Exim is typically installed as a setuid binary and most operating systems will inhibit coredumps of these by default, so further OS-specific action may be required.

<b>pipe_as_creator</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------	----------------------	-----------------------

If the generic **user** option is not set and this option is true, the delivery process is run under the uid that was in force when Exim was originally called to accept the message. If the group id is not otherwise set (via the generic **group** option), the gid that was in force when Exim was originally called to accept the message is used.

<b>restrict_to_path</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------	------------------	----------------------	-----------------------

When this option is set, any command name not listed in **allow\_commands** must contain no slashes. The command is searched for only in the directories listed in the **path** option. This option is intended for use in the case when a pipe command has been generated from a user's *.forward* file. This is usually handled by a *pipe* transport called **address\_pipe**.

<b>return_fail_output</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	------------------	----------------------	-----------------------

If this option is true, and the command produced any output and ended with a return code other than zero or one of the codes listed in **temp\_errors** (that is, the delivery failed), the output is returned in the bounce message. However, if the message has a null sender (that is, it is itself a bounce message), output from the command is discarded. This option and **return\_output** are mutually exclusive. Only one of them may be set.

<b>return_output</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------	----------------------	-----------------------

If this option is true, and the command produced any output, the delivery is deemed to have failed whatever the return code from the command, and the output is returned in the bounce message. Otherwise, the output is just discarded. However, if the message has a null sender (that is, it is a bounce message), output from the command is always discarded, whatever the setting of this option. This option and **return\_fail\_output** are mutually exclusive. Only one of them may be set.

<b>temp_errors</b>	Use: <i>pipe</i>	Type: <i>string list</i>	Default: <i>see below</i>
--------------------	------------------	--------------------------	---------------------------

This option contains either a colon-separated list of numbers, or a single asterisk. If **ignore\_status** is false and **return\_output** is not set, and the command exits with a non-zero return code, the failure is treated as temporary and the delivery is deferred if the return code matches one of the numbers, or if the setting is a single asterisk. Otherwise, non-zero return codes are treated as permanent errors. The default setting contains the codes defined by EX\_TEMPFAIL and EX\_CANTCREAT in *sysexits.h*. If Exim is compiled on a system that does not define these macros, it assumes values of 75 and 73, respectively.

<b>timeout</b>	Use: <i>pipe</i>	Type: <i>time</i>	Default: <i>1h</i>
----------------	------------------	-------------------	--------------------

If the command fails to complete within this time, it is killed. This normally causes the delivery to fail (but see **timeout\_defer**). A zero time interval specifies no timeout. In order to ensure that any subprocesses created by the command are also killed, Exim makes the initial process a process group leader, and kills the whole process group on a timeout. However, this can be defeated if one of the processes starts a new process group.

<b>timeout_defer</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------	----------------------	-----------------------

A timeout in a *pipe* transport, either in the command that the transport runs, or in a transport filter that is associated with it, is by default treated as a hard error, and the delivery fails. However, if **timeout\_defer** is set true, both kinds of timeout become temporary errors, causing the delivery to be deferred.

<b>umask</b>	Use: <i>pipe</i>	Type: <i>octal integer</i>	Default: <i>022</i>
--------------	------------------	----------------------------	---------------------

This specifies the umask setting for the subprocess that runs the command.

<b>use_bsmtplib</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------	------------------	----------------------	-----------------------

If this option is set true, the *pipe* transport writes messages in “batch SMTP” format, with the envelope sender and recipient(s) included as SMTP commands. If you want to include a leading HELO command with such messages, you can do so by setting the **message\_prefix** option. See section 48.10 for details of batch SMTP.

<b>use_classresources</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	------------------	----------------------	-----------------------

This option is available only when Exim is running on FreeBSD, NetBSD, or BSD/OS. If it is set true, the *setclassresources()* function is used to set resource limits when a *pipe* transport is run to perform a delivery. The limits for the uid under which the pipe is to run are obtained from the login class database.

<b>use_crlf</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------	------------------	----------------------	-----------------------

This option causes lines to be terminated with the two-character CRLF sequence (carriage return, linefeed) instead of just a linefeed character. In the case of batched SMTP, the byte sequence written to the pipe is then an exact image of what would be sent down a real SMTP connection.

The contents of the **message\_prefix** and **message\_suffix** options are written verbatim, so must contain their own carriage return characters if these are needed. When **use\_bsmtplib** is not set, the default values for both **message\_prefix** and **message\_suffix** end with a single linefeed, so their values must be changed to end with `\r\n` if **use\_crlf** is set.

<b>use_shell</b>	Use: <i>pipe</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------	------------------	----------------------	-----------------------

If this option is set, it causes the command to be passed to */bin/sh* instead of being run directly from the transport, as described in section 29.3. This is less secure, but is needed in some situations where the command is expected to be run under a shell and cannot easily be modified. The **allow\_commands** and **restrict\_to\_path** options, and the `$pipe_addresses` facility are incompatible with **use\_shell**. The command is expanded as a single string, and handed to */bin/sh* as data for its **-c** option.

## 29.6 Using an external local delivery agent

The *pipe* transport can be used to pass all messages that require local delivery to a separate local delivery agent such as **procmail**. When doing this, care must be taken to ensure that the pipe is run under an appropriate uid and gid. In some configurations one wants this to be a uid that is trusted by the delivery agent to supply the correct sender of the message. It may be necessary to recompile or reconfigure the delivery agent so that it trusts an appropriate user. The following is an example transport and router configuration for **procmail**:

```
# transport
procmail_pipe:
  driver = pipe
  command = /usr/local/bin/procmail -d $local_part
  return_path_add
  delivery_date_add
  envelope_to_add
  check_string = "From "
  escape_string = ">From "
  umask = 077
  user = $local_part
  group = mail

# router
procmail:
  driver = accept
  check_local_user
  transport = procmail_pipe
```

In this example, the pipe is run as the local user, but with the group set to *mail*. An alternative is to run the pipe as a specific user such as *mail* or *exim*, but in this case you must arrange for **procmail** to trust that user to supply a correct sender address. If you do not specify either a **group** or a **user** option, the pipe command is run as the local user. The home directory is the user's home directory by default.

**Note:** The command that the pipe transport runs does *not* begin with

```
IFS=" "
```

as shown in some **procmail** documentation, because Exim does not by default use a shell to run pipe commands.

The next example shows a transport and a router for a system where local deliveries are handled by the Cyrus IMAP server.

```
# transport
local_delivery_cyrus:
  driver = pipe
  command = /usr/cyrus/bin/deliver \
            -m ${substr_1:$local_part_suffix} -- $local_part
  user = cyrus
  group = mail
  return_output
  log_output
  message_prefix =
```

```
message_suffix =  
  
# router  
local_user_cyrus:  
    driver = accept  
    check_local_user  
    local_part_suffix = .*  
    transport = local_delivery_cyrus
```

Note the unsetting of **message\_prefix** and **message\_suffix**, and the use of **return\_output** to cause any text written by Cyrus to be returned to the sender.



## 30. The smtp transport

The *smtp* transport delivers messages over TCP/IP connections using the SMTP or LMTP protocol. The list of hosts to try can either be taken from the address that is being processed (having been set up by the router), or specified explicitly for the transport. Timeout and retry processing (see chapter 32) is applied to each IP address independently.

### 30.1 Multiple messages on a single connection

The sending of multiple messages over a single TCP/IP connection can arise in two ways:

- If a message contains more than **max\_rcpt** (see below) addresses that are routed to the same host, more than one copy of the message has to be sent to that host. In this situation, multiple copies may be sent in a single run of the *smtp* transport over a single TCP/IP connection. (What Exim actually does when it has too many addresses to send in one message also depends on the value of the global **remote\_max\_parallel** option. Details are given in section 48.1.)
- When a message has been successfully delivered over a TCP/IP connection, Exim looks in its hints database to see if there are any other messages awaiting a connection to the same host. If there are, a new delivery process is started for one of them, and the current TCP/IP connection is passed on to it. The new process may in turn send multiple copies and possibly create yet another process.

For each copy sent over the same TCP/IP connection, a sequence counter is incremented, and if it ever gets to the value of **connection\_max\_messages**, no further messages are sent over that connection.

### 30.2 Use of the \$host and \$host\_address variables

At the start of a run of the *smtp* transport, the values of *\$host* and *\$host\_address* are the name and IP address of the first host on the host list passed by the router. However, when the transport is about to connect to a specific host, and while it is connected to that host, *\$host* and *\$host\_address* are set to the values for that host. These are the values that are in force when the **helo\_data**, **hosts\_try\_auth**, **interface**, **serialize\_hosts**, and the various TLS options are expanded.

### 30.3 Use of \$tls\_cipher and \$tls\_peerdn

At the start of a run of the *smtp* transport, the values of *\$tls\_bits*, *\$tls\_cipher*, *\$tls\_peerdn* and *\$tls\_sni* are the values that were set when the message was received. These are the values that are used for options that are expanded before any SMTP connections are made. Just before each connection is made, these four variables are emptied. If TLS is subsequently started, they are set to the appropriate values for the outgoing connection, and these are the values that are in force when any authenticators are run and when the **authenticated\_sender** option is expanded.

These variables are deprecated in favour of *\$tls\_in\_cipher* et. al. and will be removed in a future release.

### 30.4 Private options for smtp

The private options of the *smtp* transport are as follows:

<b>address_retry_include_sender</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>true</i>
-------------------------------------	------------------	----------------------	----------------------

When an address is delayed because of a 4xx response to a RCPT command, it is the combination of sender and recipient that is delayed in subsequent queue runs until the retry time is reached. You can delay the recipient without reference to the sender (which is what earlier versions of Exim did), by setting **address\_retry\_include\_sender** false. However, this can lead to problems with servers that regularly issue 4xx responses to RCPT commands.

<b>allow_localhost</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------	----------------------	-----------------------

When a host specified in **hosts** or **fallback\_hosts** (see below) turns out to be the local host, or is listed in **hosts\_treat\_as\_local**, delivery is deferred by default. However, if **allow\_localhost** is set, Exim goes on to do the delivery anyway. This should be used only in special cases when the configuration ensures that no looping will result (for example, a differently configured Exim is listening on the port to which the message is sent).

<b>authenticated_sender</b>	Use: <i>smtp</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-----------------------------	------------------	----------------------------------	-----------------------

When Exim has authenticated as a client, or if **authenticated\_sender\_force** is true, this option sets a value for the AUTH= item on outgoing MAIL commands, overriding any existing authenticated sender value. If the string expansion is forced to fail, the option is ignored. Other expansion failures cause delivery to be deferred. If the result of expansion is an empty string, that is also ignored.

The expansion happens after the outgoing connection has been made and TLS started, if required. This means that the *\$host*, *\$host\_address*, *\$tls\_out\_cipher*, and *\$tls\_out\_peerdn* variables are set according to the particular connection.

If the SMTP session is not authenticated, the expansion of **authenticated\_sender** still happens (and can cause the delivery to be deferred if it fails), but no AUTH= item is added to MAIL commands unless **authenticated\_sender\_force** is true.

This option allows you to use the *smtp* transport in LMTP mode to deliver mail to Cyrus IMAP and provide the proper local part as the “authenticated sender”, via a setting such as:

```
authenticated_sender = $local_part
```

This removes the need for IMAP subfolders to be assigned special ACLs to allow direct delivery to those subfolders.

Because of expected uses such as that just described for Cyrus (when no domain is involved), there is no checking on the syntax of the provided value.

<b>authenticated_sender_force</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------------------	------------------	----------------------	-----------------------

If this option is set true, the **authenticated\_sender** option’s value is used for the AUTH= item on outgoing MAIL commands, even if Exim has not authenticated as a client.

<b>command_timeout</b>	Use: <i>smtp</i>	Type: <i>time</i>	Default: <i>5m</i>
------------------------	------------------	-------------------	--------------------

This sets a timeout for receiving a response to an SMTP command that has been sent out. It is also used when waiting for the initial banner line from the remote host. Its value must not be zero.

<b>connect_timeout</b>	Use: <i>smtp</i>	Type: <i>time</i>	Default: <i>5m</i>
------------------------	------------------	-------------------	--------------------

This sets a timeout for the *connect()* function, which sets up a TCP/IP call to a remote host. A setting of zero allows the system timeout (typically several minutes) to act. To have any effect, the value of this option must be less than the system timeout. However, it has been observed that on some systems there is no system timeout, which is why the default value for this option is 5 minutes, a value recommended by RFC 1123.

<b>connection_max_messages</b>	Use: <i>smtp</i>	Type: <i>integer</i>	Default: <i>500</i>
--------------------------------	------------------	----------------------	---------------------

This controls the maximum number of separate message deliveries that are sent over a single TCP/IP connection. If the value is zero, there is no limit. For testing purposes, this value can be overridden by the **-oB** command line option.

<b>data_timeout</b>	Use: <i>smtp</i>	Type: <i>time</i>	Default: <i>5m</i>
---------------------	------------------	-------------------	--------------------

This sets a timeout for the transmission of each block in the data portion of the message. As a result, the overall timeout for a message depends on the size of the message. Its value must not be zero. See also **final\_timeout**.

<b>dkim_domain</b>	Use: <i>smtp</i>	Type: <i>string</i>	Default: <i>list†</i>
--------------------	------------------	---------------------	-----------------------

<b>dkim_selector</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

<b>dkim_private_key</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	------------------	----------------------	-----------------------

<b>dkim_canon</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------	------------------	----------------------	-----------------------

<b>dkim_strict</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------	------------------	----------------------	-----------------------

<b>dkim_sign_headers</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>per RFC</i>
--------------------------	------------------	----------------------	-------------------------

<b>dkim_hash</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>sha256</i>
------------------	------------------	----------------------	------------------------

<b>dkim_identity</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

DKIM signing options. For details see section 57.1.

<b>delay_after_cutoff</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------------------	------------------	----------------------	----------------------

This option controls what happens when all remote IP addresses for a given domain have been inaccessible for so long that they have passed their retry cutoff times.

In the default state, if the next retry time has not been reached for any of them, the address is bounced without trying any deliveries. In other words, Exim delays retrying an IP address after the final cutoff time until a new retry time is reached, and can therefore bounce an address without ever trying a delivery, when machines have been down for a long time. Some people are unhappy at this prospect, so...

If **delay\_after\_cutoff** is set false, Exim behaves differently. If all IP addresses are past their final cutoff time, Exim tries to deliver to those IP addresses that have not been tried since the message arrived. If there are none, or if they all fail, the address is bounced. In other words, it does not delay when a new message arrives, but immediately tries those expired IP addresses that haven't been tried since the message arrived. If there is a continuous stream of messages for the dead hosts, unsetting **delay\_after\_cutoff** means that there will be many more attempts to deliver to them.

<b>dns_qualify_single</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>true</i>
---------------------------	------------------	----------------------	----------------------

If the **hosts** or **fallback\_hosts** option is being used, and the **gethostbyname** option is false, the RES\_DEFNAMES resolver option is set. See the **qualify\_single** option in chapter 17 for more details.

<b>dns_search_parents</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>false</i>
---------------------------	------------------	----------------------	-----------------------

If the **hosts** or **fallback\_hosts** option is being used, and the **gethostbyname** option is false, the RES\_DNSRCH resolver option is set. See the **search\_parents** option in chapter 17 for more details.

<b>dnssec_request_domains</b>	Use: <i>smtp</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
-------------------------------	------------------	---------------------------	-----------------------

DNS lookups for domains matching **dnssec\_request\_domains** will be done with the dnssec request bit set. This applies to all of the SRV, MX, AAAA, A lookup sequence.

<b>dnssec_require_domains</b>	Use: <i>smtp</i>	Type: <i>domain list†</i>	Default: <i>unset</i>
-------------------------------	------------------	---------------------------	-----------------------

DNS lookups for domains matching **dnssec\_require\_domains** will be done with the dnssec request bit set. Any returns not having the Authenticated Data bit (AD bit) set will be ignored and logged as a host-lookup failure. This applies to all of the SRV, MX, AAAA, A lookup sequence.

<b>dscp</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------	------------------	----------------------	-----------------------

This option causes the DSCP value associated with a socket to be set to one of a number of fixed strings or to numeric value. The **-bl:dscp** option may be used to ask Exim which names it knows of. Common values include *throughput*, *mincost*, and on newer systems *ef*, *af41*, etc. Numeric values may be in the range 0 to 0x3F.

The outbound packets from Exim will be marked with this value in the header (for IPv4, the TOS field; for IPv6, the TCLASS field); there is no guarantee that these values will have any effect, not be stripped by networking equipment, or do much of anything without cooperation with your Network Engineer and those of all network operators between the source and destination.

<b>fallback_hosts</b>	Use: <i>smtp</i>	Type: <i>string list</i>	Default: <i>unset</i>
-----------------------	------------------	--------------------------	-----------------------

String expansion is not applied to this option. The argument must be a colon-separated list of host names or IP addresses, optionally also including port numbers, though the separator can be changed, as described in section 6.20. Each individual item in the list is the same as an item in a **route\_list** setting for the *manualroute* router, as described in section 20.5.

Fallback hosts can also be specified on routers, which associate them with the addresses they process. As for the **hosts** option without **hosts\_override**, **fallback\_hosts** specified on the transport is used only if the address does not have its own associated fallback host list. Unlike **hosts**, a setting of **fallback\_hosts** on an address is not overridden by **hosts\_override**. However, **hosts\_randomize** does apply to fallback host lists.

If Exim is unable to deliver to any of the hosts for a particular address, and the errors are not permanent rejections, the address is put on a separate transport queue with its host list replaced by the fallback hosts, unless the address was routed via MX records and the current host was in the original MX list. In that situation, the fallback host list is not used.

Once normal deliveries are complete, the fallback queue is delivered by re-running the same transports with the new host lists. If several failing addresses have the same fallback hosts (and **max\_rcpt** permits it), a single copy of the message is sent.

The resolution of the host names on the fallback list is controlled by the **gethostbyname** option, as for the **hosts** option. Fallback hosts apply both to cases when the host list comes with the address and when it is taken from **hosts**. This option provides a “use a smart host only if delivery fails” facility.

<b>final_timeout</b>	Use: <i>smtp</i>	Type: <i>time</i>	Default: <i>10m</i>
----------------------	------------------	-------------------	---------------------

This is the timeout that applies while waiting for the response to the final line containing just “.” that terminates a message. Its value must not be zero.

<b>gethostbyname</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>false</i>
----------------------	------------------	----------------------	-----------------------

If this option is true when the **hosts** and/or **fallback\_hosts** options are being used, names are looked up using *gethostbyname()* (or *getipnodebyname()* when available) instead of using the DNS. Of course, that function may in fact use the DNS, but it may also consult other sources of information such as */etc/hosts*.

<b>gnutls_compat_mode</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>unset</i>
---------------------------	------------------	----------------------	-----------------------

This option controls whether GnuTLS is used in compatibility mode in an Exim server. This reduces security slightly, but improves interworking with older implementations of TLS.

<b>helo_data</b>	Use: <i>smtp</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
------------------	------------------	----------------------------------	---------------------------

The value of this option is expanded after a connection to a another host has been set up. The result is used as the argument for the EHLO, HELO, or LHLO command that starts the outgoing SMTP or LMTP session. The default value of the option is:

```
$primary_hostname
```

During the expansion, the variables *\$host* and *\$host\_address* are set to the identity of the remote host, and the variables *\$sending\_ip\_address* and *\$sending\_port* are set to the local IP address and port number that are being used. These variables can be used to generate different values for different servers or different local IP addresses. For example, if you want the string that is used for **helo\_data** to be obtained by a DNS lookup of the outgoing interface address, you could use this:

```
helo_data = ${lookup dnsdb{ptr=$sending_ip_address}{$value}\
  {primary_hostname}}
```

The use of **helo\_data** applies both to sending messages and when doing callouts.

<b>hosts</b>	Use: <i>smtp</i>	Type: <i>string list</i> <sup>†</sup>	Default: <i>unset</i>
--------------	------------------	---------------------------------------	-----------------------

Hosts are associated with an address by a router such as *dnslookup*, which finds the hosts by looking up the address domain in the DNS, or by *manualroute*, which has lists of hosts in its configuration. However, email addresses can be passed to the *smtp* transport by any router, and not all of them can provide an associated list of hosts.

The **hosts** option specifies a list of hosts to be used if the address being processed does not have any hosts associated with it. The hosts specified by **hosts** are also used, whether or not the address has its own hosts, if **hosts\_override** is set.

The string is first expanded, before being interpreted as a colon-separated list of host names or IP addresses, possibly including port numbers. The separator may be changed to something other than colon, as described in section 6.20. Each individual item in the list is the same as an item in a **route\_list** setting for the *manualroute* router, as described in section 20.5. However, note that the */MX* facility of the *manualroute* router is not available here.

If the expansion fails, delivery is deferred. Unless the failure was caused by the inability to complete a lookup, the error is logged to the panic log as well as the main log. Host names are looked up either by searching directly for address records in the DNS or by calling *gethostbyname()* (or *getipnodebyname()* when available), depending on the setting of the **gethostbyname** option. When Exim is compiled with IPv6 support, if a host that is looked up in the DNS has both IPv4 and IPv6 addresses, both types of address are used.

During delivery, the hosts are tried in order, subject to their retry status, unless **hosts\_randomize** is set.

<b>hosts_avoid_esmtp</b>	Use: <i>smtp</i>	Type: <i>host list</i> †	Default: <i>unset</i>
--------------------------	------------------	--------------------------	-----------------------

This option is for use with broken hosts that announce ESMTP facilities (for example, PIPELINING) and then fail to implement them properly. When a host matches **hosts\_avoid\_esmtp**, Exim sends HELO rather than EHLO at the start of the SMTP session. This means that it cannot use any of the ESMTP facilities such as AUTH, PIPELINING, SIZE, and STARTTLS.

<b>hosts_avoid_pipelining</b>	Use: <i>smtp</i>	Type: <i>host list</i> †	Default: <i>unset</i>
-------------------------------	------------------	--------------------------	-----------------------

Exim will not use the SMTP PIPELINING extension when delivering to any host that matches this list, even if the server host advertises PIPELINING support.

<b>hosts_avoid_tls</b>	Use: <i>smtp</i>	Type: <i>host list</i> †	Default: <i>unset</i>
------------------------	------------------	--------------------------	-----------------------

Exim will not try to start a TLS session when delivering to any host that matches this list. See chapter 42 for details of TLS.

<b>hosts_verify_avoid_tls</b>	Use: <i>smtp</i>	Type: <i>host list</i> †	Default: <i>unset</i>
-------------------------------	------------------	--------------------------	-----------------------

Exim will not try to start a TLS session for a verify callout, or when delivering in cutthrough mode, to any host that matches this list.

<b>hosts_max_try</b>	Use: <i>smtp</i>	Type: <i>integer</i>	Default: 5
----------------------	------------------	----------------------	------------

This option limits the number of IP addresses that are tried for any one delivery in cases where there are temporary delivery errors. Section 30.5 describes in detail how the value of this option is used.

<b>hosts_max_try_hardlimit</b>	Use: <i>smtp</i>	Type: <i>integer</i>	Default: 50
--------------------------------	------------------	----------------------	-------------

This is an additional check on the maximum number of IP addresses that Exim tries for any one delivery. Section 30.5 describes its use and why it exists.

<b>hosts_nopass_tls</b>	Use: <i>smtp</i>	Type: <i>host list</i> †	Default: <i>unset</i>
-------------------------	------------------	--------------------------	-----------------------

For any host that matches this list, a connection on which a TLS session has been started will not be passed to a new delivery process for sending another message on the same connection. See section 42.11 for an explanation of when this might be needed.

<b>hosts_noproxy_tls</b>	Use: <i>smtp</i>	Type: <i>host list</i> †	Default: *
--------------------------	------------------	--------------------------	------------

For any host that matches this list, a TLS session which has been started will not be passed to a new delivery process for sending another message on the same session.

The traditional implementation closes down TLS and re-starts it in the new process, on the same open TCP connection, for each successive message sent. If permitted by this option a pipe to the new process is set up instead, and the original process maintains the TLS connection and proxies the SMTP connection from and to the new process and any subsequents. The new process has no access to TLS information, so cannot include it in logging.

<b>hosts_override</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>false</i>
-----------------------	------------------	----------------------	-----------------------

If this option is set and the **hosts** option is also set, any hosts that are attached to the address are ignored, and instead the hosts specified by the **hosts** option are always used. This option does not apply to **fallback\_hosts**.

<b>hosts_randomize</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------	------------------	----------------------	-----------------------

If this option is set, and either the list of hosts is taken from the **hosts** or the **fallback\_hosts** option, or the hosts supplied by the router were not obtained from MX records (this includes fallback hosts from the router), and were not randomized by the router, the order of trying the hosts is randomized each time the transport runs. Randomizing the order of a host list can be used to do crude load sharing.

When **hosts\_randomize** is true, a host list may be split into groups whose order is separately randomized. This makes it possible to set up MX-like behaviour. The boundaries between groups are indicated by an item that is just + in the host list. For example:

```
hosts = host1:host2:host3::host4:host5
```

The order of the first three hosts and the order of the last two hosts is randomized for each use, but the first three always end up before the last two. If **hosts\_randomize** is not set, a + item in the list is ignored.

<b>hosts_require_auth</b>	Use: <i>smtp</i>	Type: <i>host list†</i>	Default: <i>unset</i>
---------------------------	------------------	-------------------------	-----------------------

This option provides a list of servers for which authentication must succeed before Exim will try to transfer a message. If authentication fails for servers which are not in this list, Exim tries to send unauthenticated. If authentication fails for one of these servers, delivery is deferred. This temporary error is detectable in the retry rules, so it can be turned into a hard failure if required. See also **hosts\_try\_auth**, and chapter 33 for details of authentication.

<b>hosts_request_ocsp</b>	Use: <i>smtp</i>	Type: <i>host list†</i>	Default: <i>*</i>
---------------------------	------------------	-------------------------	-------------------

Exim will request a Certificate Status on a TLS session for any host that matches this list. **tls\_verify\_certificates** should also be set for the transport.

<b>hosts_require_ocsp</b>	Use: <i>smtp</i>	Type: <i>host list†</i>	Default: <i>unset</i>
---------------------------	------------------	-------------------------	-----------------------

Exim will request, and check for a valid Certificate Status being given, on a TLS session for any host that matches this list. **tls\_verify\_certificates** should also be set for the transport.

<b>hosts_require_tls</b>	Use: <i>smtp</i>	Type: <i>host list†</i>	Default: <i>unset</i>
--------------------------	------------------	-------------------------	-----------------------

Exim will insist on using a TLS session when delivering to any host that matches this list. See chapter 42 for details of TLS. **Note:** This option affects outgoing mail only. To insist on TLS for incoming messages, use an appropriate ACL.

<b>hosts_try_auth</b>	Use: <i>smtp</i>	Type: <i>host list†</i>	Default: <i>unset</i>
-----------------------	------------------	-------------------------	-----------------------

This option provides a list of servers to which, provided they announce authentication support, Exim will attempt to authenticate as a client when it connects. If authentication fails, Exim will try to transfer the message unauthenticated. See also **hosts\_require\_auth**, and chapter 33 for details of authentication.

<b>hosts_try_chunking</b>	Use: <i>smtp</i>	Type: <i>host list†</i>	Default: <i>*</i>
---------------------------	------------------	-------------------------	-------------------

This option provides a list of servers to which, provided they announce CHUNKING support, Exim will attempt to use BDAT commands rather than DATA. BDAT will not be used in conjunction with a transport filter.

<b>hosts_try_fastopen</b>	Use: <i>smtp</i>	Type: <i>host list†</i>	Default: <i>unset</i>
---------------------------	------------------	-------------------------	-----------------------

This option provides a list of servers to which, provided the facility is supported by this system, Exim will attempt to perform a TCP Fast Open. No data is sent on the SYN segment but, if the remote server also supports the facility, it can send its SMTP banner immediately after the SYN,ACK segment. This can save up to one round-trip time.

The facility is only active for previously-contacted servers, as the initiator must present a cookie in the SYN segment.

On (at least some) current Linux distributions the facility must be enabled in the kernel by the sysadmin before the support is usable. There is no option for control of the server side; if the system supports it it is always enabled. Note that lengthy operations in the connect ACL, such as DNSBL lookups, will still delay the emission of the SMTP banner.

<b>hosts_try_prdr</b>	Use: <i>smtp</i>	Type: <i>host list†</i>	Default: <i>*</i>
-----------------------	------------------	-------------------------	-------------------

This option provides a list of servers to which, provided they announce PRDR support, Exim will attempt to negotiate PRDR for multi-recipient messages. The option can usually be left as default.

<b>interface</b>	Use: <i>smtp</i>	Type: <i>string list†</i>	Default: <i>unset</i>
------------------	------------------	---------------------------	-----------------------

This option specifies which interface to bind to when making an outgoing SMTP call. The value is an IP address, not an interface name such as `eth0`. Do not confuse this with the interface address that was used when a message was received, which is in `$received_ip_address`, formerly known as `$interface_address`. The name was changed to minimize confusion with the outgoing interface address. There is no variable that contains an outgoing interface address because, unless it is set by this option, its value is unknown.

During the expansion of the **interface** option the variables `$host` and `$host_address` refer to the host to which a connection is about to be made during the expansion of the string. Forced expansion failure, or an empty string result causes the option to be ignored. Otherwise, after expansion, the string must be a list of IP addresses, colon-separated by default, but the separator can be changed in the usual way. For example:

```
interface = <; 192.168.123.123 ; 3ffe:ffff:836f::fe86:a061
```

The first interface of the correct type (IPv4 or IPv6) is used for the outgoing connection. If none of them are the correct type, the option is ignored. If **interface** is not set, or is ignored, the system's IP functions choose which interface to use if the host has more than one.

<b>keepalive</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>true</i>
------------------	------------------	----------------------	----------------------

This option controls the setting of `SO_KEEPALIVE` on outgoing TCP/IP socket connections. When set, it causes the kernel to probe idle connections periodically, by sending packets with “old” sequence numbers. The other end of the connection should send a acknowledgment if the connection is still okay or a reset if the connection has been aborted. The reason for doing this is that it has the beneficial effect of freeing up certain types of connection that can get stuck when the remote host is disconnected without tidying up the TCP/IP call properly. The keepalive mechanism takes several hours to detect unreachable hosts.



<b>lmtp_ignore_quota</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>false</i>
--------------------------	------------------	----------------------	-----------------------

If this option is set true when the **protocol** option is set to “lmtp”, the string `IGNOREQUOTA` is added to RCPT commands, provided that the LMTP server has advertised support for IGNOREQUOTA in its response to the LHLO command.

<b>max_rcpt</b>	Use: <i>smtp</i>	Type: <i>integer</i>	Default: <i>100</i>
-----------------	------------------	----------------------	---------------------

This option limits the number of RCPT commands that are sent in a single SMTP message transaction. Each set of addresses is treated independently, and so can cause parallel connections to the same host if **remote\_max\_parallel** permits this.

<b>multi_domain</b>	Use: <i>smtp</i>	Type: <i>boolean</i> <sup>†</sup>	Default: <i>true</i>
---------------------	------------------	-----------------------------------	----------------------

When this option is set, the *smtp* transport can handle a number of addresses containing a mixture of different domains provided they all resolve to the same list of hosts. Turning the option off restricts the transport to handling only one domain at a time. This is useful if you want to use *\$domain* in an expansion for the transport, because it is set only when there is a single domain involved in a remote delivery.

It is expanded per-address and can depend on any of *\$address\_data*, *\$domain\_data*, *\$local\_part\_data*, *\$host*, *\$host\_address* and *\$host\_port*.

<b>port</b>	Use: <i>smtp</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
-------------	------------------	----------------------------------	---------------------------

This option specifies the TCP/IP port on the server to which Exim connects. **Note:** Do not confuse this with the port that was used when a message was received, which is in *\$received\_port*, formerly known as *\$interface\_port*. The name was changed to minimize confusion with the outgoing port. There is no variable that contains an outgoing port.

If the value of this option begins with a digit it is taken as a port number; otherwise it is looked up using *getservbyname()*. The default value is normally “smtp”, but if **protocol** is set to “lmtp”, the default is “lmtp”. If the expansion fails, or if a port number cannot be found, delivery is deferred.

<b>protocol</b>	Use: <i>smtp</i>	Type: <i>string</i>	Default: <i>smtp</i>
-----------------	------------------	---------------------	----------------------

If this option is set to “lmtp” instead of “smtp”, the default value for the **port** option changes to “lmtp”, and the transport operates the LMTP protocol (RFC 2033) instead of SMTP. This protocol is sometimes used for local deliveries into closed message stores. Exim also has support for running LMTP over a pipe to a local process – see chapter 28.

If this option is set to “smtps”, the default value for the **port** option changes to “smtps”, and the transport initiates TLS immediately after connecting, as an outbound SSL-on-connect, instead of using STARTTLS to upgrade. The Internet standards bodies strongly discourage use of this mode.

<b>retry_include_ip_address</b>	Use: <i>smtp</i>	Type: <i>boolean</i> <sup>†</sup>	Default: <i>true</i>
---------------------------------	------------------	-----------------------------------	----------------------

Exim normally includes both the host name and the IP address in the key it constructs for indexing retry data after a temporary delivery failure. This means that when one of several IP addresses for a host is failing, it gets tried periodically (controlled by the retry rules), but use of the other IP addresses is not affected.

However, in some dialup environments hosts are assigned a different IP address each time they connect. In this situation the use of the IP address as part of the retry key leads to undesirable behaviour. Setting this option false causes Exim to use only the host name. Since it is expanded it can be made to depend on the host or domain.

<b>serialize_hosts</b>	Use: <i>smtp</i>	Type: <i>host list</i> <sup>†</sup>	Default: <i>unset</i>
------------------------	------------------	-------------------------------------	-----------------------

Because Exim operates in a distributed manner, if several messages for the same host arrive at around the same time, more than one simultaneous connection to the remote host can occur. This is not usually a problem except when there is a slow link between the hosts. In that situation it may be helpful to restrict Exim to one connection at a time. This can be done by setting **serialize\_hosts** to match the relevant hosts.

Exim implements serialization by means of a hints database in which a record is written whenever a process connects to one of the restricted hosts. The record is deleted when the connection is completed. Obviously there is scope for records to get left lying around if there is a system or program crash. To guard against this, Exim ignores any records that are more than six hours old.

If you set up this kind of serialization, you should also arrange to delete the relevant hints database whenever your system reboots. The names of the files start with *misc* and they are kept in the *spool/db* directory. There may be one or two files, depending on the type of DBM in use. The same files are used for ETRN serialization.

See also the **max\_parallel** generic transport option.

<b>size_addition</b>	Use: <i>smtp</i>	Type: <i>integer</i>	Default: <i>1024</i>
----------------------	------------------	----------------------	----------------------

If a remote SMTP server indicates that it supports the SIZE option of the MAIL command, Exim uses this to pass over the message size at the start of an SMTP transaction. It adds the value of **size\_addition** to the value it sends, to allow for headers and other text that may be added during delivery by configuration options or in a transport filter. It may be necessary to increase this if a lot of text is added to messages.

Alternatively, if the value of **size\_addition** is set negative, it disables the use of the SIZE option altogether.

<b>socks_proxy</b>	Use: <i>smtp</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
--------------------	------------------	----------------------------------	-----------------------

This option enables use of SOCKS proxies for connections made by the transport. For details see section 58.2.

<b>tls_certificate</b>	Use: <i>smtp</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
------------------------	------------------	----------------------------------	-----------------------

The value of this option must be the absolute path to a file which contains the client's certificate, for possible use when sending a message over an encrypted connection. The values of *\$host* and *\$host\_address* are set to the name and address of the server during the expansion. See chapter 42 for details of TLS.

**Note:** This option must be set if you want Exim to be able to use a TLS certificate when sending messages as a client. The global option of the same name specifies the certificate for Exim as a server; it is not automatically assumed that the same certificate should be used when Exim is operating as a client.

<b>tls_crl</b>	Use: <i>smtp</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
----------------	------------------	----------------------------------	-----------------------

This option specifies a certificate revocation list. The expanded value must be the name of a file that contains a CRL in PEM format.

<b>tls_dh_min_bits</b>	Use: <i>smtp</i>	Type: <i>integer</i>	Default: <i>1024</i>
------------------------	------------------	----------------------	----------------------

When establishing a TLS session, if a ciphersuite which uses Diffie-Hellman key agreement is negotiated, the server will provide a large prime number for use. This option establishes the minimum

acceptable size of that number. If the parameter offered by the server is too small, then the TLS handshake will fail.

Only supported when using GnuTLS.

<b>tls_privatekey</b>	Use: <i>smtp</i>	Type: <i>string</i> †	Default: <i>unset</i>
-----------------------	------------------	-----------------------	-----------------------

The value of this option must be the absolute path to a file which contains the client's private key. This is used when sending a message over an encrypted connection using a client certificate. The values of *\$host* and *\$host\_address* are set to the name and address of the server during the expansion. If this option is unset, or the expansion is forced to fail, or the result is an empty string, the private key is assumed to be in the same file as the certificate. See chapter 42 for details of TLS.

<b>tls_require_ciphers</b>	Use: <i>smtp</i>	Type: <i>string</i> †	Default: <i>unset</i>
----------------------------	------------------	-----------------------	-----------------------

The value of this option must be a list of permitted cipher suites, for use when setting up an outgoing encrypted connection. (There is a global option of the same name for controlling incoming connections.) The values of *\$host* and *\$host\_address* are set to the name and address of the server during the expansion. See chapter 42 for details of TLS; note that this option is used in different ways by OpenSSL and GnuTLS (see sections 42.4 and 42.5). For GnuTLS, the order of the ciphers is a preference order.

<b>tls_sni</b>	Use: <i>smtp</i>	Type: <i>string</i> †	Default: <i>unset</i>
----------------	------------------	-----------------------	-----------------------

If this option is set then it sets the *\$tls\_out\_sni* variable and causes any TLS session to pass this value as the Server Name Indication extension to the remote side, which can be used by the remote side to select an appropriate certificate and private key for the session.

See 42.10 for more information.

Note that for OpenSSL, this feature requires a build of OpenSSL that supports TLS extensions.

<b>tls_tempfail_tryclear</b>	Use: <i>smtp</i>	Type: <i>boolean</i>	Default: <i>true</i>
------------------------------	------------------	----------------------	----------------------

When the server host is not in **hosts\_require\_tls**, and there is a problem in setting up a TLS session, this option determines whether or not Exim should try to deliver the message unencrypted. If it is set false, delivery to the current host is deferred; if there are other hosts, they are tried. If this option is set true, Exim attempts to deliver unencrypted after a 4xx response to STARTTLS. Also, if STARTTLS is accepted, but the subsequent TLS negotiation fails, Exim closes the current connection (because it is in an unknown state), opens a new one to the same host, and then tries the delivery in clear.

<b>tls_try_verify_hosts</b>	Use: <i>smtp</i>	Type: <i>host list</i> †	Default: <i>*</i>
-----------------------------	------------------	--------------------------	-------------------

This option gives a list of hosts for which, on encrypted connections, certificate verification will be tried but need not succeed. The **tls\_verify\_certificates** option must also be set. Note that unless the host is in this list TLS connections will be denied to hosts using self-signed certificates when **tls\_verify\_certificates** is matched. The *\$tls\_out\_certificate\_verified* variable is set when certificate verification succeeds.

<b>tls_verify_cert_hostnames</b>	Use: <i>smtp</i>	Type: <i>host list</i> †	Default: <i>*</i>
----------------------------------	------------------	--------------------------	-------------------

This option give a list of hosts for which, while verifying the server certificate, checks will be included on the host name (note that this will generally be the result of a DNS MX lookup) versus Subject and Subject-Alternate-Name fields. Wildcard names are permitted limited to being the initial component of a 3-or-more component FQDN.

There is no equivalent checking on client certificates.

<b>tls_verify_certificates</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>system</i>
--------------------------------	------------------	----------------------	------------------------

The value of this option must be either the word "system" or the absolute path to a file or directory containing permitted certificates for servers, for use when setting up an encrypted connection.

The "system" value for the option will use a location compiled into the SSL library. This is not available for GnuTLS versions preceding 3.0.20; a value of "system" is taken as empty and an explicit location must be specified.

The use of a directory for the option value is not available for GnuTLS versions preceding 3.3.6 and a single file must be used.

With OpenSSL the certificates specified explicitly either by file or directory are added to those given by the system default location.

The values of *\$host* and *\$host\_address* are set to the name and address of the server during the expansion of this option. See chapter 42 for details of TLS.

For back-compatibility, if neither `tls_verify_hosts` nor `tls_try_verify_hosts` are set (a single-colon empty list counts as being set) and certificate verification fails the TLS connection is closed.

<b>tls_verify_hosts</b>	Use: <i>smtp</i>	Type: <i>host list†</i>	Default: <i>unset</i>
-------------------------	------------------	-------------------------	-----------------------

This option gives a list of hosts for which, on encrypted connections, certificate verification must succeed. The **tls\_verify\_certificates** option must also be set. If both this option and **tls\_try\_verify\_hosts** are unset operation is as if this option selected all hosts.

## 30.5 How the limits for the number of hosts to try are used

There are two options that are concerned with the number of hosts that are tried when an SMTP delivery takes place. They are **hosts\_max\_try** and **hosts\_max\_try\_hardlimit**.

The **hosts\_max\_try** option limits the number of hosts that are tried for a single delivery. However, despite the term "host" in its name, the option actually applies to each IP address independently. In other words, a multihomed host is treated as several independent hosts, just as it is for retrying.

Many of the larger ISPs have multiple MX records which often point to multihomed hosts. As a result, a list of a dozen or more IP addresses may be created as a result of routing one of these domains.

Trying every single IP address on such a long list does not seem sensible; if several at the top of the list fail, it is reasonable to assume there is some problem that is likely to affect all of them. Roughly speaking, the value of **hosts\_max\_try** is the maximum number that are tried before deferring the delivery. However, the logic cannot be quite that simple.

Firstly, IP addresses that are skipped because their retry times have not arrived do not count, and in addition, addresses that are past their retry limits are also not counted, even when they are tried. This means that when some IP addresses are past their retry limits, more than the value of **hosts\_max\_retry** may be tried. The reason for this behaviour is to ensure that all IP addresses are considered before timing out an email address (but see below for an exception).

Secondly, when the **hosts\_max\_try** limit is reached, Exim looks down the host list to see if there is a subsequent host with a different (higher valued) MX. If there is, that host is considered next, and the current IP address is used but not counted. This behaviour helps in the case of a domain with a retry rule that hardly ever delays any hosts, as is now explained:

Consider the case of a long list of hosts with one MX value, and a few with a higher MX value. If **hosts\_max\_try** is small (the default is 5) only a few hosts at the top of the list are tried at first. With the default retry rule, which specifies increasing retry times, the higher MX hosts are eventually tried when those at the top of the list are skipped because they have not reached their retry times.

However, it is common practice to put a fixed short retry time on domains for large ISPs, on the grounds that their servers are rarely down for very long. Unfortunately, these are exactly the domains that tend to resolve to long lists of hosts. The short retry time means that the lowest MX hosts are tried every time. The attempts may be in a different order because of random sorting, but without the special MX check, the higher MX hosts would never be tried until all the lower MX hosts had timed out (which might be several days), because there are always some lower MX hosts that have reached their retry times. With the special check, Exim considers at least one IP address from each MX value at every delivery attempt, even if the **hosts\_max\_try** limit has already been reached.

The above logic means that **hosts\_max\_try** is not a hard limit, and in particular, Exim normally eventually tries all the IP addresses before timing out an email address. When **hosts\_max\_try** was implemented, this seemed a reasonable thing to do. Recently, however, some lunatic DNS configurations have been set up with hundreds of IP addresses for some domains. It can take a very long time indeed for an address to time out in these cases.

The **hosts\_max\_try\_hardlimit** option was added to help with this problem. Exim never tries more than this number of IP addresses; if it hits this limit and they are all timed out, the email address is bounced, even though not all possible IP addresses have been tried.

## 31. Address rewriting

There are some circumstances in which Exim automatically rewrites domains in addresses. The two most common are when an address is given without a domain (referred to as an “unqualified address”) or when an address contains an abbreviated domain that is expanded by DNS lookup.

Unqualified envelope addresses are accepted only for locally submitted messages, or for messages that are received from hosts matching **sender\_unqualified\_hosts** or **recipient\_unqualified\_hosts**, as appropriate. Unqualified addresses in header lines are qualified if they are in locally submitted messages, or messages from hosts that are permitted to send unqualified envelope addresses. Otherwise, unqualified addresses in header lines are neither qualified nor rewritten.

One situation in which Exim does *not* automatically rewrite a domain is when it is the name of a CNAME record in the DNS. The older RFCs suggest that such a domain should be rewritten using the “canonical” name, and some MTAs do this. The new RFCs do not contain this suggestion.

### 31.1 Explicitly configured address rewriting

This chapter describes the rewriting rules that can be used in the main rewrite section of the configuration file, and also in the generic **headers\_rewrite** option that can be set on any transport.

Some people believe that configured address rewriting is a Mortal Sin. Others believe that life is not possible without it. Exim provides the facility; you do not have to use it.

The main rewriting rules that appear in the “rewrite” section of the configuration file are applied to addresses in incoming messages, both envelope addresses and addresses in header lines. Each rule specifies the types of address to which it applies.

Whether or not addresses in header lines are rewritten depends on the origin of the headers and the type of rewriting. Global rewriting, that is, rewriting rules from the rewrite section of the configuration file, is applied only to those headers that were received with the message. Header lines that are added by ACLs or by a system filter or by individual routers or transports (which are specific to individual recipient addresses) are not rewritten by the global rules.

Rewriting at transport time, by means of the **headers\_rewrite** option, applies all headers except those added by routers and transports. That is, as well as the headers that were received with the message, it also applies to headers that were added by an ACL or a system filter.

In general, rewriting addresses from your own system or domain has some legitimacy. Rewriting other addresses should be done only with great care and in special circumstances. The author of Exim believes that rewriting should be used sparingly, and mainly for “regularizing” addresses in your own domains. Although it can sometimes be used as a routing tool, this is very strongly discouraged.

There are two commonly encountered circumstances where rewriting is used, as illustrated by these examples:

- The company whose domain is *hitch.fict.example* has a number of hosts that exchange mail with each other behind a firewall, but there is only a single gateway to the outer world. The gateway rewrites *\*.hitch.fict.example* as *hitch.fict.example* when sending mail off-site.
- A host rewrites the local parts of its own users so that, for example, *fp42@hitch.fict.example* becomes *Ford.Prefect@hitch.fict.example*.

### 31.2 When does rewriting happen?

Configured address rewriting can take place at several different stages of a message’s processing.

At the start of an ACL for MAIL, the sender address may have been rewritten by a special SMTP-time rewrite rule (see section 31.9), but no ordinary rewrite rules have yet been applied. If, however, the sender address is verified in the ACL, it is rewritten before verification, and remains rewritten thereafter. The subsequent value of *\$sender\_address* is the rewritten address. This also applies if sender verification happens in a RCPT ACL. Otherwise, when the sender address is not verified, it is rewritten as soon as a message’s header lines have been received.

Similarly, at the start of an ACL for RCPT, the current recipient's address may have been rewritten by a special SMTP-time rewrite rule, but no ordinary rewrite rules have yet been applied to it. However, the behaviour is different from the sender address when a recipient is verified. The address is rewritten for the verification, but the rewriting is not remembered at this stage. The value of *\$local\_part* and *\$domain* after verification are always the same as they were before (that is, they contain the unrewritten – except for SMTP-time rewriting – address).

As soon as a message's header lines have been received, all the envelope recipient addresses are permanently rewritten, and rewriting is also applied to the addresses in the header lines (if configured). This happens before adding any header lines that were specified in MAIL or RCPT ACLs, and before the DATA ACL and *local\_scan()* functions are run.

When an address is being routed, either for delivery or for verification, rewriting is applied immediately to child addresses that are generated by redirection, unless **no\_rewrite** is set on the router.

At transport time, additional rewriting of addresses in header lines can be specified by setting the generic **headers\_rewrite** option on a transport. This option contains rules that are identical in form to those in the rewrite section of the configuration file. They are applied to the original message header lines and any that were added by ACLs or a system filter. They are not applied to header lines that are added by routers or the transport.

The outgoing envelope sender can be rewritten by means of the **return\_path** transport option. However, it is not possible to rewrite envelope recipients at transport time.

### 31.3 Testing the rewriting rules that apply on input

Exim's input rewriting configuration appears in a part of the run time configuration file headed by "begin rewrite". It can be tested by the **-brw** command line option. This takes an address (which can be a full RFC 2822 address) as its argument. The output is a list of how the address would be transformed by the rewriting rules for each of the different places it might appear in an incoming message, that is, for each different header and for the envelope sender and recipient fields. For example,

```
exim -brw ph10@exim.workshop.example
```

might produce the output

```
sender: Philip.Hazel@exim.workshop.example
from: Philip.Hazel@exim.workshop.example
to: ph10@exim.workshop.example
cc: ph10@exim.workshop.example
bcc: ph10@exim.workshop.example
reply-to: Philip.Hazel@exim.workshop.example
env-from: Philip.Hazel@exim.workshop.example
env-to: ph10@exim.workshop.example
```

which shows that rewriting has been set up for that address when used in any of the source fields, but not when it appears as a recipient address. At the present time, there is no equivalent way of testing rewriting rules that are set for a particular transport.

### 31.4 Rewriting rules

The rewrite section of the configuration file consists of lines of rewriting rules in the form

```
<source pattern> <replacement> <flags>
```

Rewriting rules that are specified for the **headers\_rewrite** generic transport option are given as a colon-separated list. Each item in the list takes the same form as a line in the main rewriting configuration (except that any colons must be doubled, of course).

The formats of source patterns and replacement strings are described below. Each is terminated by white space, unless enclosed in double quotes, in which case normal quoting conventions apply inside

the quotes. The flags are single characters which may appear in any order. Spaces and tabs between them are ignored.

For each address that could potentially be rewritten, the rules are scanned in order, and replacements for the address from earlier rules can themselves be replaced by later rules (but see the “q” and “R” flags).

The order in which addresses are rewritten is undefined, may change between releases, and must not be relied on, with one exception: when a message is received, the envelope sender is always rewritten first, before any header lines are rewritten. For example, the replacement string for a rewrite of an address in *To*: must not assume that the message’s address in *From*: has (or has not) already been rewritten. However, a rewrite of *From*: may assume that the envelope sender has already been rewritten.

The variables *\$local\_part* and *\$domain* can be used in the replacement string to refer to the address that is being rewritten. Note that lookup-driven rewriting can be done by a rule of the form

```
*@*    ${lookup ...
```

where the lookup key uses *\$1* and *\$2* or *\$local\_part* and *\$domain* to refer to the address that is being rewritten.

## 31.5 Rewriting patterns

The source pattern in a rewriting rule is any item which may appear in an address list (see section 10.19). It is in fact processed as a single-item address list, which means that it is expanded before being tested against the address. As always, if you use a regular expression as a pattern, you must take care to escape dollar and backslash characters, or use the `\N` facility to suppress string expansion within the regular expression.

Domains in patterns should be given in lower case. Local parts in patterns are case-sensitive. If you want to do case-insensitive matching of local parts, you can use a regular expression that starts with `^(?i)`.

After matching, the numerical variables *\$1*, *\$2*, etc. may be set, depending on the type of match which occurred. These can be used in the replacement string to insert portions of the incoming address. *\$0* always refers to the complete incoming address. When a regular expression is used, the numerical variables are set from its capturing subexpressions. For other types of pattern they are set as follows:

- If a local part or domain starts with an asterisk, the numerical variables refer to the character strings matched by asterisks, with *\$1* associated with the first asterisk, and *\$2* with the second, if present. For example, if the pattern

```
*queen@*.fict.example
```

is matched against the address *hearts-queen@wonderland.fict.example* then

```
$0 = hearts-queen@wonderland.fict.example
$1 = hearts-
$2 = wonderland
```

Note that if the local part does not start with an asterisk, but the domain does, it is *\$1* that contains the wild part of the domain.

- If the domain part of the pattern is a partial lookup, the wild and fixed parts of the domain are placed in the next available numerical variables. Suppose, for example, that the address *foo@bar.baz.example* is processed by a rewriting rule of the form

```
*@partial-dbm;/some/dbm/file <replacement string>
```

and the key in the file that matches the domain is *\*.baz.example*. Then

```
$1 = foo
$2 = bar
$3 = baz.example
```



If the address *foo@baz.example* is looked up, this matches the same wildcard file entry, and in this case \$2 is set to the empty string, but \$3 is still set to *baz.example*. If a non-wild key is matched in a partial lookup, \$2 is again set to the empty string and \$3 is set to the whole domain. For non-partial domain lookups, no numerical variables are set.

## 31.6 Rewriting replacements

If the replacement string for a rule is a single asterisk, addresses that match the pattern and the flags are *not* rewritten, and no subsequent rewriting rules are scanned. For example,

```
hatta@lookingglass.fict.example * f
```

specifies that *hatta@lookingglass.fict.example* is never to be rewritten in *From:* headers.

If the replacement string is not a single asterisk, it is expanded, and must yield a fully qualified address. Within the expansion, the variables *\$local\_part* and *\$domain* refer to the address that is being rewritten. Any letters they contain retain their original case – they are not lower cased. The numerical variables are set up according to the type of pattern that matched the address, as described above. If the expansion is forced to fail by the presence of “fail” in a conditional or lookup item, rewriting by the current rule is abandoned, but subsequent rules may take effect. Any other expansion failure causes the entire rewriting operation to be abandoned, and an entry written to the panic log.

## 31.7 Rewriting flags

There are three different kinds of flag that may appear on rewriting rules:

- Flags that specify which headers and envelope addresses to rewrite: E, F, T, b, c, f, h, r, s, t.
- A flag that specifies rewriting at SMTP time: S.
- Flags that control the rewriting process: Q, q, R, w.

For rules that are part of the **headers\_rewrite** generic transport option, E, F, T, and S are not permitted.

## 31.8 Flags specifying which headers and envelope addresses to rewrite

If none of the following flag letters, nor the “S” flag (see section 31.9) are present, a main rewriting rule applies to all headers and to both the sender and recipient fields of the envelope, whereas a transport-time rewriting rule just applies to all headers. Otherwise, the rewriting rule is skipped unless the relevant addresses are being processed.

E	rewrite all envelope fields
F	rewrite the envelope From field
T	rewrite the envelope To field
b	rewrite the <i>Bcc:</i> header
c	rewrite the <i>Cc:</i> header
f	rewrite the <i>From:</i> header
h	rewrite all headers
r	rewrite the <i>Reply-To:</i> header
s	rewrite the <i>Sender:</i> header
t	rewrite the <i>To:</i> header

"All headers" means all of the headers listed above that can be selected individually, plus their *Resent-*versions. It does not include other headers such as *Subject:* etc.

You should be particularly careful about rewriting *Sender:* headers, and restrict this to special known cases in your own domains.

## 31.9 The SMTP-time rewriting flag

The rewrite flag “S” specifies a rewrite of incoming envelope addresses at SMTP time, as soon as an address is received in a MAIL or RCPT command, and before any other processing; even before

syntax checking. The pattern is required to be a regular expression, and it is matched against the whole of the data for the command, including any surrounding angle brackets.

This form of rewrite rule allows for the handling of addresses that are not compliant with RFCs 2821 and 2822 (for example, “bang paths” in batched SMTP input). Because the input is not required to be a syntactically valid address, the variables *\$local\_part* and *\$domain* are not available during the expansion of the replacement string. The result of rewriting replaces the original address in the MAIL or RCPT command.

## 31.10 Flags controlling the rewriting process

There are four flags which control the way the rewriting process works. These take effect only when a rule is invoked, that is, when the address is of the correct type (matches the flags) and matches the pattern:

- If the “Q” flag is set on a rule, the rewritten address is permitted to be an unqualified local part. It is qualified with **qualify\_recipient**. In the absence of “Q” the rewritten address must always include a domain.
- If the “q” flag is set on a rule, no further rewriting rules are considered, even if no rewriting actually takes place because of a “fail” in the expansion. The “q” flag is not effective if the address is of the wrong type (does not match the flags) or does not match the pattern.
- The “R” flag causes a successful rewriting rule to be re-applied to the new address, up to ten times. It can be combined with the “q” flag, to stop rewriting once it fails to match (after at least one successful rewrite).
- When an address in a header is rewritten, the rewriting normally applies only to the working part of the address, with any comments and RFC 2822 “phrase” left unchanged. For example, rewriting might change

```
From: Ford Prefect <fp42@restaurant.hitch.fict.example>
```

into

```
From: Ford Prefect <prefectf@hitch.fict.example>
```

Sometimes there is a need to replace the whole address item, and this can be done by adding the flag letter “w” to a rule. If this is set on a rule that causes an address in a header line to be rewritten, the entire address is replaced, not just the working part. The replacement must be a complete RFC 2822 address, including the angle brackets if necessary. If text outside angle brackets contains a character whose value is greater than 126 or less than 32 (except for tab), the text is encoded according to RFC 2047. The character set is taken from **headers\_charset**, which gets its default at build time.

When the “w” flag is set on a rule that causes an envelope address to be rewritten, all but the working part of the replacement address is discarded.

## 31.11 Rewriting examples

Here is an example of the two common rewriting paradigms:

```
*@*.hitch.fict.example    $1@hitch.fict.example
*@hitch.fict.example      ${lookup{$1}dbm{/etc/realnames}\
                           {$value}fail}@hitch.fict.example bctfrF
```

Note the use of “fail” in the lookup expansion in the second rule, forcing the string expansion to fail if the lookup does not succeed. In this context it has the effect of leaving the original address unchanged, but Exim goes on to consider subsequent rewriting rules, if any, because the “q” flag is not present in that rule. An alternative to “fail” would be to supply *\$1* explicitly, which would cause the rewritten address to be the same as before, at the cost of a small bit of processing. Not supplying either of these is an error, since the rewritten address would then contain no local part.

The first example above replaces the domain with a superior, more general domain. This may not be desirable for certain local parts. If the rule

```
root@*.hitch.fict.example *
```

were inserted before the first rule, rewriting would be suppressed for the local part *root* at any domain ending in *hitch.fict.example*.

Rewriting can be made conditional on a number of tests, by making use of *\$/if* in the expansion item. For example, to apply a rewriting rule only to messages that originate outside the local host:

```
*@*.hitch.fict.example "${if !eq {$sender_host_address}{}\
{$1@hitch.fict.example}fail}"
```

The replacement string is quoted in this example because it contains white space.

Exim does not handle addresses in the form of “bang paths”. If it sees such an address it treats it as an unqualified local part which it qualifies with the local qualification domain (if the source of the message is local or if the remote host is permitted to send unqualified addresses). Rewriting can sometimes be used to handle simple bang paths with a fixed number of components. For example, the rule

```
\N^ ([^!]+) ! (.* )@your.domain.example$ \N $2@$1
```

rewrites a two-component bang path *host.name!user* as the domain address *user@host.name*. However, there is a security implication in using this as a global rewriting rule for envelope addresses. It can provide a backdoor method for using your system as a relay, because the incoming addresses appear to be local. If the bang path addresses are received via SMTP, it is safer to use the “S” flag to rewrite them as they are received, so that relay checking can be done on the rewritten addresses.

## 32. Retry configuration

The “retry” section of the runtime configuration file contains a list of retry rules that control how often Exim tries to deliver messages that cannot be delivered at the first attempt. If there are no retry rules (the section is empty or not present), there are no retries. In this situation, temporary errors are treated as permanent. The default configuration contains a single, general-purpose retry rule (see section 7.5). The **-brt** command line option can be used to test which retry rule will be used for a given address, domain and error.

The most common cause of retries is temporary failure to deliver to a remote host because the host is down, or inaccessible because of a network problem. Exim’s retry processing in this case is applied on a per-host (strictly, per IP address) basis, not on a per-message basis. Thus, if one message has recently been delayed, delivery of a new message to the same host is not immediately tried, but waits for the host’s retry time to arrive. If the **retry\_defer** log selector is set, the message “retry time not reached” is written to the main log whenever a delivery is skipped for this reason. Section 48.2 contains more details of the handling of errors during remote deliveries.

Retry processing applies to routing as well as to delivering, except as covered in the next paragraph. The retry rules do not distinguish between these actions. It is not possible, for example, to specify different behaviour for failures to route the domain *snark.fict.example* and failures to deliver to the host *snark.fict.example*. I didn’t think anyone would ever need this added complication, so did not implement it. However, although they share the same retry rule, the actual retry times for routing and transporting a given domain are maintained independently.

When a delivery is not part of a queue run (typically an immediate delivery on receipt of a message), the routers are always run, and local deliveries are always attempted, even if retry times are set for them. This makes for better behaviour if one particular message is causing problems (for example, causing quota overflow, or provoking an error in a filter file). If such a delivery suffers a temporary failure, the retry data is updated as normal, and subsequent delivery attempts from queue runs occur only when the retry time for the local address is reached.

### 32.1 Changing retry rules

If you change the retry rules in your configuration, you should consider whether or not to delete the retry data that is stored in Exim’s spool area in files with names like *db/retry*. Deleting any of Exim’s hints files is always safe; that is why they are called “hints”.

The hints retry data contains suggested retry times based on the previous rules. In the case of a long-running problem with a remote host, it might record the fact that the host has timed out. If your new rules increase the timeout time for such a host, you should definitely remove the old retry data and let Exim recreate it, based on the new rules. Otherwise Exim might bounce messages that it should now be retaining.

### 32.2 Format of retry rules

Each retry rule occupies one line and consists of three or four parts, separated by white space: a pattern, an error name, an optional list of sender addresses, and a list of retry parameters. The pattern and sender lists must be enclosed in double quotes if they contain white space. The rules are searched in order until one is found where the pattern, error name, and sender list (if present) match the failing host or address, the error that occurred, and the message’s sender, respectively.

The pattern is any single item that may appear in an address list (see section 10.19). It is in fact processed as a one-item address list, which means that it is expanded before being tested against the address that has been delayed. A negated address list item is permitted. Address list processing treats a plain domain name as if it were preceded by “\*@”, which makes it possible for many retry rules to start with just a domain. For example,

```
lookingglass.fict.example      *   F,24h,30m;
```

provides a rule for any address in the *lookingglass.fict.example* domain, whereas

```
alice@lookingglass.fict.example * F, 24h, 30m;
```

applies only to temporary failures involving the local part **alice**. In practice, almost all rules start with a domain name pattern without a local part.

**Warning:** If you use a regular expression in a retry rule pattern, it must match a complete address, not just a domain, because that is how regular expressions work in address lists.

```
^\Nxyz\d+\.abc\.example$\N * G, 1h, 10m, 2 Wrong  
^\N[^@]+@xyz\d+\.abc\.example$\N * G, 1h, 10m, 2 Right
```

### 32.3 Choosing which retry rule to use for address errors

When Exim is looking for a retry rule after a routing attempt has failed (for example, after a DNS timeout), each line in the retry configuration is tested against the complete address only if **retry\_use\_local\_part** is set for the router. Otherwise, only the domain is used, except when matching against a regular expression, when the local part of the address is replaced with “\*”. A domain on its own can match a domain pattern, or a pattern that starts with “\*@”. By default, **retry\_use\_local\_part** is true for routers where **check\_local\_user** is true, and false for other routers.

Similarly, when Exim is looking for a retry rule after a local delivery has failed (for example, after a mailbox full error), each line in the retry configuration is tested against the complete address only if **retry\_use\_local\_part** is set for the transport (it defaults true for all local transports).

However, when Exim is looking for a retry rule after a remote delivery attempt suffers an address error (a 4xx SMTP response for a recipient address), the whole address is always used as the key when searching the retry rules. The rule that is found is used to create a retry time for the combination of the failing address and the message’s sender. It is the combination of sender and recipient that is delayed in subsequent queue runs until its retry time is reached. You can delay the recipient without regard to the sender by setting **address\_retry\_include\_sender** false in the *smtp* transport but this can lead to problems with servers that regularly issue 4xx responses to RCPT commands.

### 32.4 Choosing which retry rule to use for host and message errors

For a temporary error that is not related to an individual address (for example, a connection timeout), each line in the retry configuration is checked twice. First, the name of the remote host is used as a domain name (preceded by “\*@” when matching a regular expression). If this does not match the line, the domain from the email address is tried in a similar fashion. For example, suppose the MX records for *a.b.c.example* are

```
a.b.c.example  MX  5  x.y.z.example  
                MX  6  p.q.r.example  
                MX  7  m.n.o.example
```

and the retry rules are

```
p.q.r.example * F, 24h, 30m;  
a.b.c.example * F, 4d, 45m;
```

and a delivery to the host *x.y.z.example* suffers a connection failure. The first rule matches neither the host nor the domain, so Exim looks at the second rule. This does not match the host, but it does match the domain, so it is used to calculate the retry time for the host *x.y.z.example*. Meanwhile, Exim tries to deliver to *p.q.r.example*. If this also suffers a host error, the first retry rule is used, because it matches the host.

In other words, temporary failures to deliver to host *p.q.r.example* use the first rule to determine retry times, but for all the other hosts for the domain *a.b.c.example*, the second rule is used. The second rule is also used if routing to *a.b.c.example* suffers a temporary failure.

**Note:** The host name is used when matching the patterns, not its IP address. However, if a message is routed directly to an IP address without the use of a host name, for example, if a *manualroute* router contains a setting such as:

```
route_list = *.a.example 192.168.34.23
```

then the “host name” that is used when searching for a retry rule is the textual form of the IP address.

## 32.5 Retry rules for specific errors

The second field in a retry rule is the name of a particular error, or an asterisk, which matches any error. The errors that can be tested for are:

### **auth\_failed**

Authentication failed when trying to send to a host in the **hosts\_require\_auth** list in an *smtp* transport.

### **data\_4xx**

A 4xx error was received for an outgoing DATA command, either immediately after the command, or after sending the message’s data.

### **mail\_4xx**

A 4xx error was received for an outgoing MAIL command.

### **rcpt\_4xx**

A 4xx error was received for an outgoing RCPT command.

For the three 4xx errors, either the first or both of the x’s can be given as specific digits, for example: *mail\_45x* or *rcpt\_436*. For example, to recognize 452 errors given to RCPT commands for addresses in a certain domain, and have retries every ten minutes with a one-hour timeout, you could set up a retry rule of this form:

```
the.domain.name rcpt_452 F,1h,10m
```

These errors apply to both outgoing SMTP (the *smtp* transport) and outgoing LMTP (either the *lmtp* transport, or the *smtp* transport in LMTP mode).

### **lost\_connection**

A server unexpectedly closed the SMTP connection. There may, of course, legitimate reasons for this (host died, network died), but if it repeats a lot for the same host, it indicates something odd.

### **lookup**

A DNS lookup for a host failed. Note that a **dnslookup** router will need to have matched its **fail\_defer\_domains** option for this retry type to be usable. Also note that a **manualroute** router will probably need its **host\_find\_failed** option set to **defer**.

### **refused\_MX**

A connection to a host obtained from an MX record was refused.

### **refused\_A**

A connection to a host not obtained from an MX record was refused.

### **refused**

A connection was refused.

### **timeout\_connect\_MX**

A connection attempt to a host obtained from an MX record timed out.

### **timeout\_connect\_A**

A connection attempt to a host not obtained from an MX record timed out.

### **timeout\_connect**

A connection attempt timed out.

### **timeout\_MX**

There was a timeout while connecting or during an SMTP session with a host obtained from an MX record.

### **timeout\_A**

There was a timeout while connecting or during an SMTP session with a host not obtained from an MX record.

### timeout

There was a timeout while connecting or during an SMTP session.

### tls\_required

The server was required to use TLS (it matched **hosts\_require\_tls** in the *smtp* transport), but either did not offer TLS, or it responded with 4xx to STARTTLS, or there was a problem setting up the TLS connection.

### quota

A mailbox quota was exceeded in a local delivery by the *appendfile* transport.

### quota\_<time>

A mailbox quota was exceeded in a local delivery by the *appendfile* transport, and the mailbox has not been accessed for <time>. For example, *quota\_4d* applies to a quota error when the mailbox has not been accessed for four days.

The idea of **quota\_<time>** is to make it possible to have shorter timeouts when the mailbox is full and is not being read by its owner. Ideally, it should be based on the last time that the user accessed the mailbox. However, it is not always possible to determine this. Exim uses the following heuristic rules:

- If the mailbox is a single file, the time of last access (the “atime”) is used. As no new messages are being delivered (because the mailbox is over quota), Exim does not access the file, so this is the time of last user access.
- For a maildir delivery, the time of last modification of the *new* subdirectory is used. As the mailbox is over quota, no new files are created in the *new* subdirectory, because no new messages are being delivered. Any change to the *new* subdirectory is therefore assumed to be the result of an MUA moving a new message to the *cur* directory when it is first read. The time that is used is therefore the last time that the user read a new message.
- For other kinds of multi-file mailbox, the time of last access cannot be obtained, so a retry rule that uses this type of error field is never matched.

The quota errors apply both to system-enforced quotas and to Exim’s own quota mechanism in the *appendfile* transport. The *quota* error also applies when a local delivery is deferred because a partition is full (the ENOSPC error).

## 32.6 Retry rules for specified senders

You can specify retry rules that apply only when the failing message has a specific sender. In particular, this can be used to define retry rules that apply only to bounce messages. The third item in a retry rule can be of this form:

```
senders=<address list>
```

The retry timings themselves are then the fourth item. For example:

```
*    rcpt_4xx    senders=:    F, 1h, 30m
```

matches recipient 4xx errors for bounce messages sent to any address at any host. If the address list contains white space, it must be enclosed in quotes. For example:

```
a.domain rcpt_452 senders="xb.dom : yc.dom" G, 8h, 10m, 1.5
```

**Warning:** This facility can be unhelpful if it is used for host errors (which do not depend on the recipient). The reason is that the sender is used only to match the retry rule. Once the rule has been found for a host error, its contents are used to set a retry time for the host, and this will apply to all messages, not just those with specific senders.

When testing retry rules using **-brt**, you can supply a sender using the **-f** command line option, like this:

```
exim -f "" -brt user@dom.ain
```

If you do not set **-f** with **-brt**, a retry rule that contains a senders list is never matched.

## 32.7 Retry parameters

The third (or fourth, if a senders list is present) field in a retry rule is a sequence of retry parameter sets, separated by semicolons. Each set consists of

*<letter>,<cutoff time>,<arguments>*

The letter identifies the algorithm for computing a new retry time; the cutoff time is the time beyond which this algorithm no longer applies, and the arguments vary the algorithm's action. The cutoff time is measured from the time that the first failure for the domain (combined with the local part if relevant) was detected, not from the time the message was received.

The available algorithms are:

- *F*: retry at fixed intervals. There is a single time parameter specifying the interval.
- *G*: retry at geometrically increasing intervals. The first argument specifies a starting value for the interval, and the second a multiplier, which is used to increase the size of the interval at each retry.
- *H*: retry at randomized intervals. The arguments are as for *G*. For each retry, the previous interval is multiplied by the factor in order to get a maximum for the next interval. The minimum interval is the first argument of the parameter, and an actual interval is chosen randomly between them. Such a rule has been found to be helpful in cluster configurations when all the members of the cluster restart at once, and may therefore synchronize their queue processing times.

When computing the next retry time, the algorithm definitions are scanned in order until one whose cutoff time has not yet passed is reached. This is then used to compute a new retry time that is later than the current time. In the case of fixed interval retries, this simply means adding the interval to the current time. For geometrically increasing intervals, retry intervals are computed from the rule's parameters until one that is greater than the previous interval is found. The main configuration variable **retry\_interval\_max** limits the maximum interval between retries. It cannot be set greater than 24h, which is its default value.

A single remote domain may have a number of hosts associated with it, and each host may have more than one IP address. Retry algorithms are selected on the basis of the domain name, but are applied to each IP address independently. If, for example, a host has two IP addresses and one is unusable, Exim will generate retry times for it and will not try to use it until its next retry time comes. Thus the good IP address is likely to be tried first most of the time.

Retry times are hints rather than promises. Exim does not make any attempt to run deliveries exactly at the computed times. Instead, a queue runner process starts delivery processes for delayed messages periodically, and these attempt new deliveries only for those addresses that have passed their next retry time. If a new message arrives for a deferred address, an immediate delivery attempt occurs only if the address has passed its retry time. In the absence of new messages, the minimum time between retries is the interval between queue runner processes. There is not much point in setting retry times of five minutes if your queue runners happen only once an hour, unless there are a significant number of incoming messages (which might be the case on a system that is sending everything to a smart host, for example).

The data in the retry hints database can be inspected by using the *exim\_dumpdb* or *exim\_fixdb* utility programs (see chapter 53). The latter utility can also be used to change the data. The *exinext* utility script can be used to find out what the next retry times are for the hosts associated with a particular mail domain, and also for local deliveries that have been deferred.

## 32.8 Retry rule examples

Here are some example retry rules:

```
alice@wonderland.fict.example quota_5d F,7d,3h
wonderland.fict.example      quota_5d
wonderland.fict.example      *          F,1h,15m; G,2d,1h,2;
lookingglass.fict.example    *          F,24h,30m;
*                            refused_A  F,2h,20m;
*                            *          F,2h,15m; G,16h,1h,1.5; F,5d,8h
```



The first rule sets up special handling for mail to *alice@wonderland.fict.example* when there is an over-quota error and the mailbox has not been read for at least 5 days. Retries continue every three hours for 7 days. The second rule handles over-quota errors for all other local parts at *wonderland.fict.example*; the absence of a local part has the same effect as supplying “\*@”. As no retry algorithms are supplied, messages that fail are bounced immediately if the mailbox has not been read for at least 5 days.

The third rule handles all other errors at *wonderland.fict.example*; retries happen every 15 minutes for an hour, then with geometrically increasing intervals until two days have passed since a delivery first failed. After the first hour there is a delay of one hour, then two hours, then four hours, and so on (this is a rather extreme example).

The fourth rule controls retries for the domain *lookingglass.fict.example*. They happen every 30 minutes for 24 hours only. The remaining two rules handle all other domains, with special action for connection refusal from hosts that were not obtained from an MX record.

The final rule in a retry configuration should always have asterisks in the first two fields so as to provide a general catch-all for any addresses that do not have their own special handling. This example tries every 15 minutes for 2 hours, then with intervals starting at one hour and increasing by a factor of 1.5 up to 16 hours, then every 8 hours up to 5 days.

## 32.9 Timeout of retry data

Exim timestamps the data that it writes to its retry hints database. When it consults the data during a delivery it ignores any that is older than the value set in **retry\_data\_expire** (default 7 days). If, for example, a host hasn't been tried for 7 days, Exim will try to deliver to it immediately a message arrives, and if that fails, it will calculate a retry time as if it were failing for the first time.

This improves the behaviour for messages routed to rarely-used hosts such as MX backups. If such a host was down at one time, and happens to be down again when Exim tries a month later, using the old retry data would imply that it had been down all the time, which is not a justified assumption.

If a host really is permanently dead, this behaviour causes a burst of retries every now and again, but only if messages routed to it are rare. If there is a message at least once every 7 days the retry data never expires.

## 32.10 Long-term failures

Special processing happens when an email address has been failing for so long that the cutoff time for the last algorithm is reached. For example, using the default retry rule:

```
* * F, 2h, 15m; G, 16h, 1h, 1.5; F, 4d, 6h
```

the cutoff time is four days. Reaching the retry cutoff is independent of how long any specific message has been failing; it is the length of continuous failure for the recipient address that counts.

When the cutoff time is reached for a local delivery, or for all the IP addresses associated with a remote delivery, a subsequent delivery failure causes Exim to give up on the address, and a bounce message is generated. In order to cater for new messages that use the failing address, a next retry time is still computed from the final algorithm, and is used as follows:

For local deliveries, one delivery attempt is always made for any subsequent messages. If this delivery fails, the address fails immediately. The post-cutoff retry time is not used.

If the delivery is remote, there are two possibilities, controlled by the **delay\_after\_cutoff** option of the *smtp* transport. The option is true by default. Until the post-cutoff retry time for one of the IP addresses is reached, the failing email address is bounced immediately, without a delivery attempt taking place. After that time, one new delivery attempt is made to those IP addresses that are past their retry times, and if that still fails, the address is bounced and new retry times are computed.

In other words, when all the hosts for a given email address have been failing for a long time, Exim bounces rather than defers until one of the hosts' retry times is reached. Then it tries once, and

bounces if that attempt fails. This behaviour ensures that few resources are wasted in repeatedly trying to deliver to a broken destination, but if the host does recover, Exim will eventually notice.

If **delay\_after\_cutoff** is set false, Exim behaves differently. If all IP addresses are past their final cutoff time, Exim tries to deliver to those IP addresses that have not been tried since the message arrived. If there are no suitable IP addresses, or if they all fail, the address is bounced. In other words, it does not delay when a new message arrives, but tries the expired addresses immediately, unless they have been tried since the message arrived. If there is a continuous stream of messages for the failing domains, setting **delay\_after\_cutoff** false means that there will be many more attempts to deliver to permanently failing IP addresses than when **delay\_after\_cutoff** is true.

## 32.11 Deliveries that work intermittently

Some additional logic is needed to cope with cases where a host is intermittently available, or when a message has some attribute that prevents its delivery when others to the same address get through. In this situation, because some messages are successfully delivered, the “retry clock” for the host or address keeps getting reset by the successful deliveries, and so failing messages remain on the queue for ever because the cutoff time is never reached.

Two exceptional actions are applied to prevent this happening. The first applies to errors that are related to a message rather than a remote host. Section 48.2 has a discussion of the different kinds of error; examples of message-related errors are 4xx responses to MAIL or DATA commands, and quota failures. For this type of error, if a message’s arrival time is earlier than the “first failed” time for the error, the earlier time is used when scanning the retry rules to decide when to try next and when to time out the address.

The exceptional second action applies in all cases. If a message has been on the queue for longer than the cutoff time of any applicable retry rule for a given address, a delivery is attempted for that address, even if it is not yet time, and if this delivery fails, the address is timed out. A new retry time is not computed in this case, so that other messages for the same address are considered immediately.

## 33. SMTP authentication

The “authenticators” section of Exim’s run time configuration is concerned with SMTP authentication. This facility is an extension to the SMTP protocol, described in RFC 2554, which allows a client SMTP host to authenticate itself to a server. This is a common way for a server to recognize clients that are permitted to use it as a relay. SMTP authentication is not of relevance to the transfer of mail between servers that have no managerial connection with each other.

Very briefly, the way SMTP authentication works is as follows:

- The server advertises a number of authentication *mechanisms* in response to the client’s EHLO command.
- The client issues an AUTH command, naming a specific mechanism. The command may, optionally, contain some authentication data.
- The server may issue one or more *challenges*, to which the client must send appropriate responses. In simple authentication mechanisms, the challenges are just prompts for user names and passwords. The server does not have to issue any challenges – in some mechanisms the relevant data may all be transmitted with the AUTH command.
- The server either accepts or denies authentication.
- If authentication succeeds, the client may optionally make use of the AUTH option on the MAIL command to pass an authenticated sender in subsequent mail transactions. Authentication lasts for the remainder of the SMTP connection.
- If authentication fails, the client may give up, or it may try a different authentication mechanism, or it may try transferring mail over the unauthenticated connection.

If you are setting up a client, and want to know which authentication mechanisms the server supports, you can use Telnet to connect to port 25 (the SMTP port) on the server, and issue an EHLO command. The response to this includes the list of supported mechanisms. For example:

```
$ telnet server.example 25
Trying 192.168.34.25...
Connected to server.example.
Escape character is '^]'.
220 server.example ESMTP Exim 4.20 ...
ehlo client.example
250-server.example Hello client.example [10.8.4.5]
250-SIZE 52428800
250-PIPELINING
250-AUTH PLAIN
250 HELP
```

The second-last line of this example output shows that the server supports authentication using the PLAIN mechanism. In Exim, the different authentication mechanisms are configured by specifying *authenticator* drivers. Like the routers and transports, which authenticators are included in the binary is controlled by build-time definitions. The following are currently available, included by setting

```
AUTH_CRAM_MD5=yes
AUTH_CYRUS_SASL=yes
AUTH_DOVECOT=yes
AUTH_GSASL=yes
AUTH_HEIMDAL_GSSAPI=yes
AUTH_PLAINTEXT=yes
AUTH_SPA=yes
AUTH_TLS=yes
```

in *Local/Makefile*, respectively. The first of these supports the CRAM-MD5 authentication mechanism (RFC 2195), and the second provides an interface to the Cyrus SASL authentication library. The third is an interface to Dovecot’s authentication system, delegating the work via a socket interface.

The fourth provides an interface to the GNU SASL authentication library, which provides mechanisms but typically not data sources. The fifth provides direct access to Heimdal GSSAPI, geared for Kerberos, but supporting setting a server keytab. The sixth can be configured to support the PLAIN authentication mechanism (RFC 2595) or the LOGIN mechanism, which is not formally documented, but used by several MUAs. The seventh authenticator supports Microsoft's *Secure Password Authentication* mechanism. The eighth is an Exim authenticator but not an SMTP one; instead it can use information from a TLS negotiation.

The authenticators are configured using the same syntax as other drivers (see section 6.23). If no authenticators are required, no authentication section need be present in the configuration file. Each authenticator can in principle have both server and client functions. When Exim is receiving SMTP mail, it is acting as a server; when it is sending out messages over SMTP, it is acting as a client. Authenticator configuration options are provided for use in both these circumstances.

To make it clear which options apply to which situation, the prefixes **server\_** and **client\_** are used on option names that are specific to either the server or the client function, respectively. Server and client functions are disabled if none of their options are set. If an authenticator is to be used for both server and client functions, a single definition, using both sets of options, is required. For example:

```
cram:
  driver = cram_md5
  public_name = CRAM-MD5
  server_secret = ${if eq{$auth1}{ph10}{secret1}fail}
  client_name = ph10
  client_secret = secret2
```

The **server\_** option is used when Exim is acting as a server, and the **client\_** options when it is acting as a client.

Descriptions of the individual authenticators are given in subsequent chapters. The remainder of this chapter covers the generic options for the authenticators, followed by general discussion of the way authentication works in Exim.

**Beware:** the meaning of *\$auth1*, *\$auth2*, ... varies on a per-driver and per-mechanism basis. Please read carefully to determine which variables hold account labels such as usercodes and which hold passwords or other authenticating data.

Note that some mechanisms support two different identifiers for accounts: the *authentication id* and the *authorization id*. The contractions *authn* and *authz* are commonly encountered. The American spelling is standard here. Conceptually, authentication data such as passwords are tied to the identifier used to authenticate; servers may have rules to permit one user to act as a second user, so that after login the session is treated as though that second user had logged in. That second user is the *authorization id*. A robust configuration might confirm that the *authz* field is empty or matches the *authn* field. Often this is just ignored. The *authn* can be considered as verified data, the *authz* as an unverified request which the server might choose to honour.

A *realm* is a text string, typically a domain name, presented by a server to a client to help it select an account and credentials to use. In some mechanisms, the client and server provably agree on the realm, but clients typically can not treat the realm as secure data to be blindly trusted.

### 33.1 Generic options for authenticators

<b>client_condition</b>	Use: <i>authenticators</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
-------------------------	----------------------------	----------------------------------	-----------------------

When Exim is authenticating as a client, it skips any authenticator whose **client\_condition** expansion yields “0”, “no”, or “false”. This can be used, for example, to skip plain text authenticators when the connection is not encrypted by a setting such as:

```
client_condition = ${if !eq{$tls_out_cipher}{}}
```

<b>client_set_id</b>	Use: <i>authenticators</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	----------------------------	----------------------	-----------------------

When client authentication succeeds, this condition is expanded; the result is used in the log lines for outbound messages. Typically it will be the user name used for authentication.

<b>driver</b>	Use: <i>authenticators</i>	Type: <i>string</i>	Default: <i>unset</i>
---------------	----------------------------	---------------------	-----------------------

This option must always be set. It specifies which of the available authenticators is to be used.

<b>public_name</b>	Use: <i>authenticators</i>	Type: <i>string</i>	Default: <i>unset</i>
--------------------	----------------------------	---------------------	-----------------------

This option specifies the name of the authentication mechanism that the driver implements, and by which it is known to the outside world. These names should contain only upper case letters, digits, underscores, and hyphens (RFC 2222), but Exim in fact matches them caselessly. If **public\_name** is not set, it defaults to the driver's instance name.

<b>server_advertise_condition</b>	Use: <i>authenticators</i>	Type: <i>string†</i>	Default: <i>unset</i>
-----------------------------------	----------------------------	----------------------	-----------------------

When a server is about to advertise an authentication mechanism, the condition is expanded. If it yields the empty string, "0", "no", or "false", the mechanism is not advertised. If the expansion fails, the mechanism is not advertised. If the failure was not forced, and was not caused by a lookup defer, the incident is logged. See section 33.3 below for further discussion.

<b>server_condition</b>	Use: <i>authenticators</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	----------------------------	----------------------	-----------------------

This option must be set for a **plaintext** server authenticator, where it is used directly to control authentication. See section 34.2 for details.

For the *gssasl* authenticator, this option is required for various mechanisms; see chapter 38 for details.

For the other authenticators, **server\_condition** can be used as an additional authentication or authorization mechanism that is applied after the other authenticator conditions succeed. If it is set, it is expanded when the authenticator would otherwise return a success code. If the expansion is forced to fail, authentication fails. Any other expansion failure causes a temporary error code to be returned. If the result of a successful expansion is an empty string, "0", "no", or "false", authentication fails. If the result of the expansion is "1", "yes", or "true", authentication succeeds. For any other result, a temporary error code is returned, with the expanded string as the error text.

<b>server_debug_print</b>	Use: <i>authenticators</i>	Type: <i>string†</i>	Default: <i>unset</i>
---------------------------	----------------------------	----------------------	-----------------------

If this option is set and authentication debugging is enabled (see the **-d** command line option), the string is expanded and included in the debugging output when the authenticator is run as a server. This can help with checking out the values of variables. If expansion of the string fails, the error message is written to the debugging output, and Exim carries on processing.

<b>server_set_id</b>	Use: <i>authenticators</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	----------------------------	----------------------	-----------------------

When an Exim server successfully authenticates a client, this string is expanded using data from the authentication, and preserved for any incoming messages in the variable *\$authenticated\_id*. It is also included in the log lines for incoming messages. For example, a user/password authenticator configuration might preserve the user name that was used to authenticate, and refer to it subsequently during delivery of the message. If expansion fails, the option is ignored.

This option allows a server to discard authenticated sender addresses supplied as part of MAIL commands in SMTP connections that are authenticated by the driver on which **server\_mail\_auth\_condition** is set. The option is not used as part of the authentication process; instead its (unexpanded) value is remembered for later use. How it is used is described in the following section.

## 33.2 The AUTH parameter on MAIL commands

When a client supplied an AUTH= item on a MAIL command, Exim applies the following checks before accepting it as the authenticated sender of the message:

- If the connection is not using extended SMTP (that is, HELO was used rather than EHLO), the use of AUTH= is a syntax error.
- If the value of the AUTH= parameter is "<>", it is ignored.
- If **acl\_smtp\_mailauth** is defined, the ACL it specifies is run. While it is running, the value of *\$authenticated\_sender* is set to the value obtained from the AUTH= parameter. If the ACL does not yield "accept", the value of *\$authenticated\_sender* is deleted. The **acl\_smtp\_mailauth** ACL may not return "drop" or "discard". If it defers, a temporary error code (451) is given for the MAIL command.
- If **acl\_smtp\_mailauth** is not defined, the value of the AUTH= parameter is accepted and placed in *\$authenticated\_sender* only if the client has authenticated.
- If the AUTH= value was accepted by either of the two previous rules, and the client has authenticated, and the authenticator has a setting for the **server\_mail\_auth\_condition**, the condition is checked at this point. The value that was saved from the authenticator is expanded. If the expansion fails, or yields an empty string, "0", "no", or "false", the value of *\$authenticated\_sender* is deleted. If the expansion yields any other value, the value of *\$authenticated\_sender* is retained and passed on with the message.

When *\$authenticated\_sender* is set for a message, it is passed on to other hosts to which Exim authenticates as a client. Do not confuse this value with *\$authenticated\_id*, which is a string obtained from the authentication process, and which is not usually a complete email address.

Whenever an AUTH= value is ignored, the incident is logged. The ACL for MAIL, if defined, is run after AUTH= is accepted or ignored. It can therefore make use of *\$authenticated\_sender*. The converse is not true: the value of *\$sender\_address* is not yet set up when the **acl\_smtp\_mailauth** ACL is run.

## 33.3 Authentication on an Exim server

When Exim receives an EHLO command, it advertises the public names of those authenticators that are configured as servers, subject to the following conditions:

- The client host must match **auth\_advertise\_hosts** (default \*).
- If the **server\_advertise\_condition** option is set, its expansion must not yield the empty string, "0", "no", or "false".

The order in which the authenticators are defined controls the order in which the mechanisms are advertised.

Some mail clients (for example, some versions of Netscape) require the user to provide a name and password for authentication whenever AUTH is advertised, even though authentication may not in fact be needed (for example, Exim may be set up to allow unconditional relaying from the client by an IP address check). You can make such clients more friendly by not advertising AUTH to them. For example, if clients on the 10.9.8.0/24 network are permitted (by the ACL that runs for RCPT) to relay without authentication, you should set

```
auth_advertise_hosts = ! 10.9.8.0/24
```

so that no authentication mechanisms are advertised to them.

The **server\_advertise\_condition** controls the advertisement of individual authentication mechanisms. For example, it can be used to restrict the advertisement of a particular mechanism to encrypted connections, by a setting such as:

```
server_advertise_condition = ${if eq{$tls_in_cipher}{}{no}{yes}}
```

If the session is encrypted, *\$tls\_in\_cipher* is not empty, and so the expansion yields “yes”, which allows the advertisement to happen.

When an Exim server receives an AUTH command from a client, it rejects it immediately if AUTH was not advertised in response to an earlier EHLO command. This is the case if

- The client host does not match **auth\_advertise\_hosts**; or
- No authenticators are configured with server options; or
- Expansion of **server\_advertise\_condition** blocked the advertising of all the server authenticators.

Otherwise, Exim runs the ACL specified by **acl\_smtp\_auth** in order to decide whether to accept the command. If **acl\_smtp\_auth** is not set, AUTH is accepted from any client host.

If AUTH is not rejected by the ACL, Exim searches its configuration for a server authentication mechanism that was advertised in response to EHLO and that matches the one named in the AUTH command. If it finds one, it runs the appropriate authentication protocol, and authentication either succeeds or fails. If there is no matching advertised mechanism, the AUTH command is rejected with a 504 error.

When a message is received from an authenticated host, the value of *\$received\_protocol* is set to “esmtpa” or “esmtpsa” instead of “esmtpt” or “esmtp”, and *\$sender\_host\_authenticated* contains the name (not the public name) of the authenticator driver that successfully authenticated the client from which the message was received. This variable is empty if there was no successful authentication.

## 33.4 Testing server authentication

Exim’s **-bh** option can be useful for testing server authentication configurations. The data for the AUTH command has to be sent using base64 encoding. A quick way to produce such data for testing is the following Perl script:

```
use MIME::Base64;
printf ("%s", encode_base64(eval "\"$ARGV[0]\""));
```

This interprets its argument as a Perl string, and then encodes it. The interpretation as a Perl string allows binary zeros, which are required for some kinds of authentication, to be included in the data. For example, a command line to run this script on such data might be

```
encode '\0user\0password'
```

Note the use of single quotes to prevent the shell interpreting the backslashes, so that they can be interpreted by Perl to specify characters whose code value is zero.

**Warning 1:** If either of the user or password strings starts with an octal digit, you must use three zeros instead of one after the leading backslash. If you do not, the octal digit that starts your string will be incorrectly interpreted as part of the code for the first character.

**Warning 2:** If there are characters in the strings that Perl interprets specially, you must use a Perl escape to prevent them being misinterpreted. For example, a command such as

```
encode '\0user@domain.com\0pas$$word'
```

gives an incorrect answer because of the unescaped “@” and “\$” characters.

If you have the **mimencode** command installed, another way to do produce base64-encoded strings is to run the command

```
echo -e -n '\0user\0password' | mimencode
```

The **-e** option of **echo** enables the interpretation of backslash escapes in the argument, and the **-n** option specifies no newline at the end of its output. However, not all versions of **echo** recognize these options, so you should check your version before relying on this suggestion.

### 33.5 Authentication by an Exim client

The *smtp* transport has two options called **hosts\_require\_auth** and **hosts\_try\_auth**. When the *smtp* transport connects to a server that announces support for authentication, and the host matches an entry in either of these options, Exim (as a client) tries to authenticate as follows:

- For each authenticator that is configured as a client, in the order in which they are defined in the configuration, it searches the authentication mechanisms announced by the server for one whose name matches the public name of the authenticator.
- When it finds one that matches, it runs the authenticator's client code. The variables *\$host* and *\$host\_address* are available for any string expansions that the client might do. They are set to the server's name and IP address. If any expansion is forced to fail, the authentication attempt is abandoned, and Exim moves on to the next authenticator. Otherwise an expansion failure causes delivery to be deferred.
- If the result of the authentication attempt is a temporary error or a timeout, Exim abandons trying to send the message to the host for the moment. It will try again later. If there are any backup hosts available, they are tried in the usual way.
- If the response to authentication is a permanent error (5xx code), Exim carries on searching the list of authenticators and tries another one if possible. If all authentication attempts give permanent errors, or if there are no attempts because no mechanisms match (or option expansions force failure), what happens depends on whether the host matches **hosts\_require\_auth** or **hosts\_try\_auth**. In the first case, a temporary error is generated, and delivery is deferred. The error can be detected in the retry rules, and thereby turned into a permanent error if you wish. In the second case, Exim tries to deliver the message unauthenticated.

Note that the hostlist test for whether to do authentication can be confused if name-IP lookups change between the time the peer is decided on and the transport running. For example, with a manualroute router given a host name, and DNS "round-robin" use by that name: if the local resolver cache times out between the router and the transport running, the transport may get an IP for the name for its authentication check which does not match the connection peer IP. No authentication will then be done, despite the names being identical.

For such cases use a separate transport which always authenticates.

When Exim has authenticated itself to a remote server, it adds the AUTH parameter to the MAIL commands it sends, if it has an authenticated sender for the message. If the message came from a remote host, the authenticated sender is the one that was receiving on an incoming MAIL command, provided that the incoming connection was authenticated and the **server\_mail\_auth** condition allowed the authenticated sender to be retained. If a local process calls Exim to send a message, the sender address that is built from the login name and **qualify\_domain** is treated as authenticated. However, if the **authenticated\_sender** option is set on the *smtp* transport, it overrides the authenticated sender that was received with the message.



## 34. The plaintext authenticator

The *plaintext* authenticator can be configured to support the PLAIN and LOGIN authentication mechanisms, both of which transfer authentication data as plain (unencrypted) text (though base64 encoded). The use of plain text is a security risk; you are strongly advised to insist on the use of SMTP encryption (see chapter 42) if you use the PLAIN or LOGIN mechanisms. If you do use unencrypted plain text, you should not use the same passwords for SMTP connections as you do for login accounts.

### 34.1 Plaintext options

When configured as a server, *plaintext* uses the following options:

<b>server_condition</b>	Use: <i>authenticators</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	----------------------------	----------------------	-----------------------

This is actually a global authentication option, but it must be set in order to configure the *plaintext* driver as a server. Its use is described below.

<b>server_prompts</b>	Use: <i>plaintext</i>	Type: <i>string†</i>	Default: <i>unset</i>
-----------------------	-----------------------	----------------------	-----------------------

The contents of this option, after expansion, must be a colon-separated list of prompt strings. If expansion fails, a temporary authentication rejection is given.

### 34.2 Using plaintext in a server

When running as a server, *plaintext* performs the authentication test by expanding a string. The data sent by the client with the AUTH command, or in response to subsequent prompts, is base64 encoded, and so may contain any byte values when decoded. If any data is supplied with the command, it is treated as a list of strings, separated by NULs (binary zeros), the first three of which are placed in the expansion variables *\$auth1*, *\$auth2*, and *\$auth3* (neither LOGIN nor PLAIN uses more than three strings).

For compatibility with previous releases of Exim, the values are also placed in the expansion variables *\$1*, *\$2*, and *\$3*. However, the use of these variables for this purpose is now deprecated, as it can lead to confusion in string expansions that also use them for other things.

If there are more strings in **server\_prompts** than the number of strings supplied with the AUTH command, the remaining prompts are used to obtain more data. Each response from the client may be a list of NUL-separated strings.

Once a sufficient number of data strings have been received, **server\_condition** is expanded. If the expansion is forced to fail, authentication fails. Any other expansion failure causes a temporary error code to be returned. If the result of a successful expansion is an empty string, “0”, “no”, or “false”, authentication fails. If the result of the expansion is “1”, “yes”, or “true”, authentication succeeds and the generic **server\_set\_id** option is expanded and saved in *\$authenticated\_id*. For any other result, a temporary error code is returned, with the expanded string as the error text

**Warning:** If you use a lookup in the expansion to find the user’s password, be sure to make the authentication fail if the user is unknown. There are good and bad examples at the end of the next section.

### 34.3 The PLAIN authentication mechanism

The PLAIN authentication mechanism (RFC 2595) specifies that three strings be sent as one item of data (that is, one combined string containing two NUL separators). The data is sent either as part of the AUTH command, or subsequently in response to an empty prompt from the server.

The second and third strings are a user name and a corresponding password. Using a single fixed user name and password as an example, this could be configured as follows:

```
fixed_plain:
    driver = plaintext
    public_name = PLAIN
    server_prompts = :
    server_condition = \
        ${if and {{eq{$auth2}{username}}{eq{$auth3}{mysecret}}}}
    server_set_id = $auth2
```

Note that the default result strings from **if** (“true” or an empty string) are exactly what we want here, so they need not be specified. Obviously, if the password contains expansion-significant characters such as dollar, backslash, or closing brace, they have to be escaped.

The **server\_prompts** setting specifies a single, empty prompt (empty items at the end of a string list are ignored). If all the data comes as part of the AUTH command, as is commonly the case, the prompt is not used. This authenticator is advertised in the response to EHLO as

```
250-AUTH PLAIN
```

and a client host can authenticate itself by sending the command

```
AUTH PLAIN AHVzZXJuYW11AG15c2VjcmV0
```

As this contains three strings (more than the number of prompts), no further data is required from the client. Alternatively, the client may just send

```
AUTH PLAIN
```

to initiate authentication, in which case the server replies with an empty prompt. The client must respond with the combined data string.

The data string is base64 encoded, as required by the RFC. This example, when decoded, is `<NUL>username<NUL>mysecret`, where `<NUL>` represents a zero byte. This is split up into three strings, the first of which is empty. The **server\_condition** option in the authenticator checks that the second two are username and mysecret respectively.

Having just one fixed user name and password, as in this example, is not very realistic, though for a small organization with only a handful of authenticating clients it could make sense.

A more sophisticated instance of this authenticator could use the user name in `$auth2` to look up a password in a file or database, and maybe do an encrypted comparison (see **crypteq** in chapter 11). Here is an example of this approach, where the passwords are looked up in a DBM file. **Warning:** This is an incorrect example:

```
server_condition = \
    ${if eq{$auth3}{{lookup{$auth2}dbm{/etc/authpwd}}}}
```

The expansion uses the user name (`$auth2`) as the key to look up a password, which it then compares to the supplied password (`$auth3`). Why is this example incorrect? It works fine for existing users, but consider what happens if a non-existent user name is given. The lookup fails, but as no success/failure strings are given for the lookup, it yields an empty string. Thus, to defeat the authentication, all a client has to do is to supply a non-existent user name and an empty password. The correct way of writing this test is:

```
server_condition = ${lookup{$auth2}dbm{/etc/authpwd}\
    ${if eq{$value}{$auth3}}} {false}}
```

In this case, if the lookup succeeds, the result is checked; if the lookup fails, “false” is returned and authentication fails. If **crypteq** is being used instead of **eq**, the first example is in fact safe, because **crypteq** always fails if its second argument is empty. However, the second way of writing the test makes the logic clearer.

## 34.4 The LOGIN authentication mechanism

The LOGIN authentication mechanism is not documented in any RFC, but is in use in a number of programs. No data is sent with the AUTH command. Instead, a user name and password are supplied

separately, in response to prompts. The plaintext authenticator can be configured to support this as in this example:

```
fixed_login:
  driver = plaintext
  public_name = LOGIN
  server_prompts = User Name : Password
  server_condition = \
    ${if and {{eq{$auth1}{username}}{eq{$auth2}{mysecret}}}}
  server_set_id = $auth1
```

Because of the way plaintext operates, this authenticator accepts data supplied with the AUTH command (in contravention of the specification of LOGIN), but if the client does not supply it (as is the case for LOGIN clients), the prompt strings are used to obtain two data items.

Some clients are very particular about the precise text of the prompts. For example, Outlook Express is reported to recognize only “Username:” and “Password:”. Here is an example of a LOGIN authenticator that uses those strings. It uses the **ldapauth** expansion condition to check the user name and password by binding to an LDAP server:

```
login:
  driver = plaintext
  public_name = LOGIN
  server_prompts = Username:: : Password::
  server_condition = ${if and{{ \
    !eq{{}{$auth1} }}{ \
    ldapauth{\
      user="uid=${quote_ldap_dn:$auth1},ou=people,o=example.org" \
      pass=${quote:$auth2} \
      ldap://ldap.example.org/} }} }
  server_set_id = uid=$auth1,ou=people,o=example.org
```

We have to check that the username is not empty before using it, because LDAP does not permit empty DN components. We must also use the **quote\_ldap\_dn** operator to correctly quote the DN for authentication. However, the basic **quote** operator, rather than any of the LDAP quoting operators, is the correct one to use for the password, because quoting is needed only to make the password conform to the Exim syntax. At the LDAP level, the password is an uninterpreted string.

## 34.5 Support for different kinds of authentication

A number of string expansion features are provided for the purpose of interfacing to different ways of user authentication. These include checking traditionally encrypted passwords from */etc/passwd* (or equivalent), PAM, Radius, **ldapauth**, *pwcheck*, and *saslauthd*. For details see section 11.7.

## 34.6 Using plaintext in a client

The *plaintext* authenticator has two client options:

<b>client_ignore_invalid_base64</b>	Use: <i>plaintext</i>	Type: <i>boolean</i>	Default: <i>false</i>
-------------------------------------	-----------------------	----------------------	-----------------------

If the client receives a server prompt that is not a valid base64 string, authentication is abandoned by default. However, if this option is set true, the error in the challenge is ignored and the client sends the response as usual.

<b>client_send</b>	Use: <i>plaintext</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
--------------------	-----------------------	----------------------------------	-----------------------

The string is a colon-separated list of authentication data strings. Each string is independently expanded before being sent to the server. The first string is sent with the AUTH command; any more strings are sent in response to prompts from the server. Before each string is expanded, the value of

the most recent prompt is placed in the next `$auth<n>` variable, starting with `$auth1` for the first prompt. Up to three prompts are stored in this way. Thus, the prompt that is received in response to sending the first string (with the AUTH command) can be used in the expansion of the second string, and so on. If an invalid base64 string is received when `client_ignore_invalid_base64` is set, an empty string is put in the `$auth<n>` variable.

**Note:** You cannot use expansion to create multiple strings, because splitting takes priority and happens first.

Because the PLAIN authentication mechanism requires NUL (binary zero) bytes in the data, further processing is applied to each string before it is sent. If there are any single circumflex characters in the string, they are converted to NULs. Should an actual circumflex be required as data, it must be doubled in the string.

This is an example of a client configuration that implements the PLAIN authentication mechanism with a fixed user name and password:

```
fixed_plain:
  driver = plaintext
  public_name = PLAIN
  client_send = ^username^mysecret
```

The lack of colons means that the entire text is sent with the AUTH command, with the circumflex characters converted to NULs. A similar example that uses the LOGIN mechanism is:

```
fixed_login:
  driver = plaintext
  public_name = LOGIN
  client_send = : username : mysecret
```

The initial colon means that the first string is empty, so no data is sent with the AUTH command itself. The remaining strings are sent in response to prompts.

## 35. The cram\_md5 authenticator

The CRAM-MD5 authentication mechanism is described in RFC 2195. The server sends a challenge string to the client, and the response consists of a user name and the CRAM-MD5 digest of the challenge string combined with a secret string (password) which is known to both server and client. Thus, the secret is not sent over the network as plain text, which makes this authenticator more secure than *plaintext*. However, the downside is that the secret has to be available in plain text at either end.

### 35.1 Using cram\_md5 as a server

This authenticator has one server option, which must be set to configure the authenticator as a server:

<b>server_secret</b>	Use: <i>cram_md5</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	----------------------	----------------------	-----------------------

When the server receives the client's response, the user name is placed in the expansion variable *\$auth1*, and **server\_secret** is expanded to obtain the password for that user. The server then computes the CRAM-MD5 digest that the client should have sent, and checks that it received the correct string. If the expansion of **server\_secret** is forced to fail, authentication fails. If the expansion fails for some other reason, a temporary error code is returned to the client.

For compatibility with previous releases of Exim, the user name is also placed in *\$1*. However, the use of this variables for this purpose is now deprecated, as it can lead to confusion in string expansions that also use numeric variables for other things.

For example, the following authenticator checks that the user name given by the client is “ph10”, and if so, uses “secret” as the password. For any other user name, authentication fails.

```
fixed_cram:
  driver = cram_md5
  public_name = CRAM-MD5
  server_secret = ${if eq{$auth1}{ph10}{secret}fail}
  server_set_id = $auth1
```

If authentication succeeds, the setting of **server\_set\_id** preserves the user name in *\$authenticated\_id*. A more typical configuration might look up the secret string in a file, using the user name as the key. For example:

```
lookup_cram:
  driver = cram_md5
  public_name = CRAM-MD5
  server_secret = ${lookup{$auth1}lsearch{/etc/authpwd}\
                  {$value}fail}
  server_set_id = $auth1
```

Note that this expansion explicitly forces failure if the lookup fails because *\$auth1* contains an unknown user name.

As another example, if you wish to re-use a Cyrus SASL *sasldb2* file without using the relevant libraries, you need to know the realm to specify in the lookup and then ask for the “userPassword” attribute for that user in that realm, with:

```
cyrusless_crammd5:
  driver = cram_md5
  public_name = CRAM-MD5
  server_secret = ${lookup{$auth1:mail.example.org:userPassword}\
                  dbmjlz{/etc/sasldb2}{$value}fail}
  server_set_id = $auth1
```

### 35.2 Using cram\_md5 as a client

When used as a client, the *cram\_md5* authenticator has two options:

<b>client_name</b>	Use: <i>cram_md5</i>	Type: <i>string†</i>	Default: <i>the primary host name</i>
--------------------	----------------------	----------------------	---------------------------------------

This string is expanded, and the result used as the user name data when computing the response to the server's challenge.

<b>client_secret</b>	Use: <i>cram_md5</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	----------------------	----------------------	-----------------------

This option must be set for the authenticator to work as a client. Its value is expanded and the result used as the secret string when computing the response.

Different user names and secrets can be used for different servers by referring to *\$host* or *\$host\_address* in the options. Forced failure of either expansion string is treated as an indication that this authenticator is not prepared to handle this case. Exim moves on to the next configured client authenticator. Any other expansion failure causes Exim to give up trying to send the message to the current server.

A simple example configuration of a *cram\_md5* authenticator, using fixed strings, is:

```
fixed_cram:
  driver = cram_md5
  public_name = CRAM-MD5
  client_name = ph10
  client_secret = secret
```

## 36. The cyrus\_sasl authenticator

The code for this authenticator was provided by Matthew Byng-Maddick of A L Digital Ltd (<http://www.aldigital.co.uk>).

The *cyrus\_sasl* authenticator provides server support for the Cyrus SASL library implementation of the RFC 2222 (“Simple Authentication and Security Layer”). This library supports a number of authentication mechanisms, including PLAIN and LOGIN, but also several others that Exim does not support directly. In particular, there is support for Kerberos authentication.

The *cyrus\_sasl* authenticator provides a gatewaying mechanism directly to the Cyrus interface, so if your Cyrus library can do, for example, CRAM-MD5, then so can the *cyrus\_sasl* authenticator. By default it uses the public name of the driver to determine which mechanism to support.

Where access to some kind of secret file is required, for example in GSSAPI or CRAM-MD5, it is worth noting that the authenticator runs as the Exim user, and that the Cyrus SASL library has no way of escalating privileges by default. You may also find you need to set environment variables, depending on the driver you are using.

The application name provided by Exim is “exim”, so various SASL options may be set in *exim.conf* in your SASL directory. If you are using GSSAPI for Kerberos, note that because of limitations in the GSSAPI interface, changing the server keytab might need to be communicated down to the Kerberos layer independently. The mechanism for doing so is dependent upon the Kerberos implementation.

For example, for older releases of Heimdal, the environment variable KRB5\_KTNAME may be set to point to an alternative keytab file. Exim will pass this variable through from its own inherited environment when started as root or the Exim user. The keytab file needs to be readable by the Exim user. With newer releases of Heimdal, a setuid Exim may cause Heimdal to discard the environment variable. In practice, for those releases, the Cyrus authenticator is not a suitable interface for GSSAPI (Kerberos) support. Instead, consider the *heimdal\_gssapi* authenticator, described in chapter 39

### 36.1 Using cyrus\_sasl as a server

The *cyrus\_sasl* authenticator has four private options. It puts the username (on a successful authentication) into *\$auth1*. For compatibility with previous releases of Exim, the username is also placed in *\$1*. However, the use of this variable for this purpose is now deprecated, as it can lead to confusion in string expansions that also use numeric variables for other things.

<b>server_hostname</b>	Use: <i>cyrus_sasl</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
------------------------	------------------------	----------------------------------	---------------------------

This option selects the hostname that is used when communicating with the library. The default value is *\$primary\_hostname*. It is up to the underlying SASL plug-in what it does with this data.

<b>server_mech</b>	Use: <i>cyrus_sasl</i>	Type: <i>string</i>	Default: <i>see below</i>
--------------------	------------------------	---------------------	---------------------------

This option selects the authentication mechanism this driver should use. The default is the value of the generic **public\_name** option. This option allows you to use a different underlying mechanism from the advertised name. For example:

```
sasl:
  driver = cyrus_sasl
  public_name = X-ANYTHING
  server_mech = CRAM-MD5
  server_set_id = $auth1
```

<b>server_realm</b>	Use: <i>cyrus_sasl</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
---------------------	------------------------	----------------------------------	-----------------------

This specifies the SASL realm that the server claims to be in.

<b>server_service</b>	Use: <i>cyrus_sasl</i>	Type: <i>string</i>	Default: <i>smtp</i>
-----------------------	------------------------	---------------------	----------------------

This is the SASL service that the server claims to implement.

For straightforward cases, you do not need to set any of the authenticator's private options. All you need to do is to specify an appropriate mechanism as the public name. Thus, if you have a SASL library that supports CRAM-MD5 and PLAIN, you could have two authenticators as follows:

```
sasl_cram_md5:
    driver = cyrus_sasl
    public_name = CRAM-MD5
    server_set_id = $auth1

sasl_plain:
    driver = cyrus_sasl
    public_name = PLAIN
    server_set_id = $auth2
```

Cyrus SASL does implement the LOGIN authentication method, even though it is not a standard method. It is disabled by default in the source distribution, but it is present in many binary distributions.



## 37. The dovecot authenticator

This authenticator is an interface to the authentication facility of the Dovecot POP/IMAP server, which can support a number of authentication methods. Note that Dovecot must be configured to use `auth-client` not `auth-userdb`. If you are using Dovecot to authenticate POP/IMAP clients, it might be helpful to use the same mechanisms for SMTP authentication. This is a server authenticator only. There is only one option:

<b>server_socket</b>	Use: <i>dovecot</i>	Type: <i>string</i>	Default: <i>unset</i>
----------------------	---------------------	---------------------	-----------------------

This option must specify the socket that is the interface to Dovecot authentication. The **public\_name** option must specify an authentication mechanism that Dovecot is configured to support. You can have several authenticators for different mechanisms. For example:

```
dovecot_plain:
  driver = dovecot
  public_name = PLAIN
  server_socket = /var/run/dovecot/auth-client
  server_set_id = $auth1

dovecot_ntlm:
  driver = dovecot
  public_name = NTLM
  server_socket = /var/run/dovecot/auth-client
  server_set_id = $auth1
```

If the SMTP connection is encrypted, or if `$sender_host_address` is equal to `$received_ip_address` (that is, the connection is local), the “secured” option is passed in the Dovecot authentication command. If, for a TLS connection, a client certificate has been verified, the “valid-client-cert” option is passed. When authentication succeeds, the identity of the user who authenticated is placed in `$auth1`.

## 38. The gssasl authenticator

The *gssasl* authenticator provides server integration for the GNU SASL library and the mechanisms it provides. This is new as of the 4.80 release and there are a few areas where the library does not let Exim smoothly scale to handle future authentication mechanisms, so no guarantee can be made that any particular new authentication mechanism will be supported without code changes in Exim.

<b>server_channelbinding</b>	Use: <i>gssasl</i>	Type: <i>boolean</i>	Default: <i>false</i>
------------------------------	--------------------	----------------------	-----------------------

Some authentication mechanisms are able to use external context at both ends of the session to bind the authentication to that context, and fail the authentication process if that context differs. Specifically, some TLS ciphersuites can provide identifying information about the cryptographic context.

This means that certificate identity and verification becomes a non-issue, as a man-in-the-middle attack will cause the correct client and server to see different identifiers and authentication will fail.

This is currently only supported when using the GnuTLS library. This is only usable by mechanisms which support "channel binding"; at time of writing, that's the SCRAM family.

This defaults off to ensure smooth upgrade across Exim releases, in case this option causes some clients to start failing. Some future release of Exim may switch the default to be true.

<b>server_hostname</b>	Use: <i>gssasl</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>see below</i>
------------------------	--------------------	----------------------------------	---------------------------

This option selects the hostname that is used when communicating with the library. The default value is `$primary_hostname`. Some mechanisms will use this data.

<b>server_mech</b>	Use: <i>gssasl</i>	Type: <i>string</i>	Default: <i>see below</i>
--------------------	--------------------	---------------------	---------------------------

This option selects the authentication mechanism this driver should use. The default is the value of the generic **public\_name** option. This option allows you to use a different underlying mechanism from the advertised name. For example:

```
sasl:
  driver = gssasl
  public_name = X-ANYTHING
  server_mech = CRAM-MD5
  server_set_id = $auth1
```

<b>server_password</b>	Use: <i>gssasl</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
------------------------	--------------------	----------------------------------	-----------------------

Various mechanisms need access to the cleartext password on the server, so that proof-of-possession can be demonstrated on the wire, without sending the password itself.

The data available for lookup varies per mechanism. In all cases, `$auth1` is set to the *authentication id*. The `$auth2` variable will always be the *authorization id (authz)* if available, else the empty string. The `$auth3` variable will always be the *realm* if available, else the empty string.

A forced failure will cause authentication to defer.

If using this option, it may make sense to set the **server\_condition** option to be simply "true".

<b>server_realm</b>	Use: <i>gssasl</i>	Type: <i>string</i> <sup>†</sup>	Default: <i>unset</i>
---------------------	--------------------	----------------------------------	-----------------------

This specifies the SASL realm that the server claims to be in. Some mechanisms will use this data.

<b>server_scam_iter</b>	Use: <i>gsasl</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	-------------------	----------------------	-----------------------

This option provides data for the SCRAM family of mechanisms. *\$auth1* is not available at evaluation time. (This may change, as we receive feedback on use)

<b>server_scam_salt</b>	Use: <i>gsasl</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------------	-------------------	----------------------	-----------------------

This option provides data for the SCRAM family of mechanisms. *\$auth1* is not available at evaluation time. (This may change, as we receive feedback on use)

<b>server_service</b>	Use: <i>gsasl</i>	Type: <i>string</i>	Default: <i>smtp</i>
-----------------------	-------------------	---------------------	----------------------

This is the SASL service that the server claims to implement. Some mechanisms will use this data.

## 38.1 gsasl auth variables

These may be set when evaluating specific options, as detailed above. They will also be set when evaluating **server\_condition**.

Unless otherwise stated below, the *gsasl* integration will use the following meanings for these variables:

- *\$auth1*: the *authentication id*
- *\$auth2*: the *authorization id*
- *\$auth3*: the *realm*

On a per-mechanism basis:

- **EXTERNAL**: only *\$auth1* is set, to the possibly empty *authorization id*; the **server\_condition** option must be present.
- **ANONYMOUS**: only *\$auth1* is set, to the possibly empty *anonymous token*; the **server\_condition** option must be present.
- **GSSAPI**: *\$auth1* will be set to the *GSSAPI Display Name*; *\$auth2* will be set to the *authorization id*, the **server\_condition** option must be present.

An *anonymous token* is something passed along as an unauthenticated identifier; this is analogous to FTP anonymous authentication passing an email address, or software-identifier@, as the "password".

An example showing the password having the realm specified in the callback and demonstrating a Cyrus SASL to GSASL migration approach is:

```
gsasl_cyrusless_crammd5:
  driver = gsasl
  public_name = CRAM-MD5
  server_realm = imap.example.org
  server_password = ${lookup{$auth1:$auth3:userPassword}\
                    dbm:jz{/etc/sasl2db2}{$value}fail}
  server_set_id = ${quote:$auth1}
  server_condition = yes
```

## 39. The heimdal\_gssapi authenticator

The *heimdal\_gssapi* authenticator provides server integration for the Heimdal GSSAPI/Kerberos library, permitting Exim to set a keytab pathname reliably.

<b>server_hostname</b>	Use: <i>heimdal_gssapi</i>	Type: <i>string†</i>	Default: <i>see below</i>
------------------------	----------------------------	----------------------	---------------------------

This option selects the hostname that is used, with **server\_service**, for constructing the GSS server name, as a *GSS\_C\_NT\_HOSTBASED\_SERVICE* identifier. The default value is *\$primary\_hostname*.

<b>server_keytab</b>	Use: <i>heimdal_gssapi</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	----------------------------	----------------------	-----------------------

If set, then Heimdal will not use the system default keytab (typically */etc/krb5.keytab*) but instead the pathname given in this option. The value should be a pathname, with no “file:” prefix.

<b>server_service</b>	Use: <i>heimdal_gssapi</i>	Type: <i>string†</i>	Default: <i>smtp</i>
-----------------------	----------------------------	----------------------	----------------------

This option specifies the service identifier used, in conjunction with **server\_hostname**, for building the identifier for finding credentials from the keytab.

### 39.1 heimdal\_gssapi auth variables

Beware that these variables will typically include a realm, thus will appear to be roughly like an email address already. The *authzid* in *\$auth2* is not verified, so a malicious client can set it to anything.

The *\$auth1* field should be safely trustable as a value from the Key Distribution Center. Note that these are not quite email addresses. Each identifier is for a role, and so the left-hand-side may include a role suffix. For instance, “joe/admin@EXAMPLE.ORG”.

- *\$auth1*: the *authentication id*, set to the GSS Display Name.
- *\$auth2*: the *authorization id*, sent within SASL encapsulation after authentication. If that was empty, this will also be set to the GSS Display Name.

## 40. The spa authenticator

The *spa* authenticator provides client support for Microsoft's *Secure Password Authentication* mechanism, which is also sometimes known as NTLM (NT LanMan). The code for client side of this authenticator was contributed by Marc Prud'hommeaux, and much of it is taken from the Samba project (<http://www.samba.org>). The code for the server side was subsequently contributed by Tom Kistner. The mechanism works as follows:

- After the AUTH command has been accepted, the client sends an SPA authentication request based on the user name and optional domain.
- The server sends back a challenge.
- The client builds a challenge response which makes use of the user's password and sends it to the server, which then accepts or rejects it.

Encryption is used to protect the password in transit.

### 40.1 Using spa as a server

The *spa* authenticator has just one server option:

<b>server_password</b>	Use: <i>spa</i>	Type: <i>string†</i>	Default: <i>unset</i>
------------------------	-----------------	----------------------	-----------------------

This option is expanded, and the result must be the cleartext password for the authenticating user, whose name is at this point in *\$auth1*. For compatibility with previous releases of Exim, the user name is also placed in *\$1*. However, the use of this variable for this purpose is now deprecated, as it can lead to confusion in string expansions that also use numeric variables for other things. For example:

```
spa:
  driver = spa
  public_name = NTLM
  server_password = \
    ${lookup{$auth1}lsearch{/etc/exim/spa_clearpass}{$value}fail}
```

If the expansion is forced to fail, authentication fails. Any other expansion failure causes a temporary error code to be returned.

### 40.2 Using spa as a client

The *spa* authenticator has the following client options:

<b>client_domain</b>	Use: <i>spa</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	-----------------	----------------------	-----------------------

This option specifies an optional domain for the authentication.

<b>client_password</b>	Use: <i>spa</i>	Type: <i>string†</i>	Default: <i>unset</i>
------------------------	-----------------	----------------------	-----------------------

This option specifies the user's password, and must be set.

<b>client_username</b>	Use: <i>spa</i>	Type: <i>string†</i>	Default: <i>unset</i>
------------------------	-----------------	----------------------	-----------------------

This option specifies the user name, and must be set. Here is an example of a configuration of this authenticator for use with the mail servers at *msn.com*:

```
msn:
  driver = spa
```

```
public_name = MSN  
client_username = msn/msn_username  
client_password = msn_plaintext_password  
client_domain = DOMAIN_OR_UNSET
```

## 41. The `tls` authenticator

The `tls` authenticator provides server support for authentication based on client certificates.

It is not an SMTP authentication mechanism and is not advertised by the server as part of the SMTP EHLO response. It is an Exim authenticator in the sense that it affects the protocol element of the log line, can be tested for by the **authenticated** ACL condition, and can set the `$authenticated_id` variable.

The client must present a verifiable certificate, for which it must have been requested via the **tls\_verify\_hosts** or **tls\_try\_verify\_hosts** main options (see 42).

If an authenticator of this type is configured it is run before any SMTP-level communication is done, and can authenticate the connection. If it does, SMTP authentication is not offered.

A maximum of one authenticator of this type may be present.

The `tls` authenticator has three server options:

<b>server_param1</b>	Use: <i>tls</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	-----------------	----------------------	-----------------------

This option is expanded after the TLS negotiation and the result is placed in `$auth1`. If the expansion is forced to fail, authentication fails. Any other expansion failure causes a temporary error code to be returned.

<b>server_param2</b>	Use: <i>tls</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	-----------------	----------------------	-----------------------

<b>server_param3</b>	Use: <i>tls</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	-----------------	----------------------	-----------------------

As above, for `$auth2` and `$auth3`.

**server\_param1** may also be spelled **server\_param**.

Example:

```
tls:
  driver = tls
  server_param1 =      ${certextract {subj_altname,mail,>:} \
                        {${tls_in_peercert}}}
  server_condition =  ${if forany {${auth1}} \
                        {!= {0} \
                        {${lookup ldap{ldap:/// \
                        mailname=${quote_ldap_dn:${lc:$item}}, \
                        ou=users,LDAP_DC?mailid} {$value}{0} \
                        } } } }
  server_set_id =      ${if = {1}}{${listcount:${auth1}} {${auth1}}{}}
```

This accepts a client certificate that is verifiable against any of your configured trust-anchors (which usually means the full set of public CAs) and which has a SAN with a good account name. Note that the client cert is on the wire in-clear, including the SAN, whereas a plaintext SMTP AUTH done inside TLS is not.

Note that because authentication is traditionally an SMTP operation, the **authenticated** ACL condition cannot be used in a connect- or helo-ACL.

## 42. Encrypted SMTP connections using TLS/SSL

Support for TLS (Transport Layer Security), formerly known as SSL (Secure Sockets Layer), is implemented by making use of the OpenSSL library or the GnuTLS library (Exim requires GnuTLS release 1.0 or later). There is no cryptographic code in the Exim distribution itself for implementing TLS. In order to use this feature you must install OpenSSL or GnuTLS, and then build a version of Exim that includes TLS support (see section 4.7). You also need to understand the basic concepts of encryption at a managerial level, and in particular, the way that public keys, private keys, and certificates are used.

RFC 3207 defines how SMTP connections can make use of encryption. Once a connection is established, the client issues a STARTTLS command. If the server accepts this, the client and the server negotiate an encryption mechanism. If the negotiation succeeds, the data that subsequently passes between them is encrypted.

Exim's ACLs can detect whether the current SMTP session is encrypted or not, and if so, what cipher suite is in use, whether the client supplied a certificate, and whether or not that certificate was verified. This makes it possible for an Exim server to deny or accept certain commands based on the encryption state.

**Warning:** Certain types of firewall and certain anti-virus products can disrupt TLS connections. You need to turn off SMTP scanning for these products in order to get TLS to work.

### 42.1 Support for the legacy “ssmtp” (aka “smtps”) protocol

Early implementations of encrypted SMTP used a different TCP port from normal SMTP, and expected an encryption negotiation to start immediately, instead of waiting for a STARTTLS command from the client using the standard SMTP port. The protocol was called “ssmtp” or “smtps”, and port 465 was allocated for this purpose.

This approach was abandoned when encrypted SMTP was standardized, but there are still some legacy clients that use it. Exim supports these clients by means of the **tls\_on\_connect\_ports** global option. Its value must be a list of port numbers; the most common use is expected to be:

```
tls_on_connect_ports = 465
```

The port numbers specified by this option apply to all SMTP connections, both via the daemon and via *inetd*. You still need to specify all the ports that the daemon uses (by setting **daemon\_smtp\_ports** or **local\_interfaces** or the **-oX** command line option) because **tls\_on\_connect\_ports** does not add an extra port – rather, it specifies different behaviour on a port that is defined elsewhere.

There is also a **-tls-on-connect** command line option. This overrides **tls\_on\_connect\_ports**; it forces the legacy behaviour for all ports.

### 42.2 OpenSSL vs GnuTLS

The first TLS support in Exim was implemented using OpenSSL. Support for GnuTLS followed later, when the first versions of GnuTLS were released. To build Exim to use GnuTLS, you need to set

```
USE_GNUTLS=yes
```

in Local/Makefile, in addition to

```
SUPPORT_TLS=yes
```

You must also set **TLS\_LIBS** and **TLS\_INCLUDE** appropriately, so that the include files and libraries for GnuTLS can be found.

There are some differences in usage when using GnuTLS instead of OpenSSL:

- The **tls\_verify\_certificates** option cannot be the path of a directory for GnuTLS versions before 3.3.6 (for later versions, or OpenSSL, it can be either).
- The default value for **tls\_dhparam** differs for historical reasons.



- Distinguished Name (DN) strings reported by the OpenSSL library use a slash for separating fields; GnuTLS uses commas, in accordance with RFC 2253. This affects the value of the `$tls_in_peerdn` and `$tls_out_peerdn` variables.
- OpenSSL identifies cipher suites using hyphens as separators, for example: DES-CBC3-SHA. GnuTLS historically used underscores, for example: RSA\_ARCFOUR\_SHA. What is more, OpenSSL complains if underscores are present in a cipher list. To make life simpler, Exim changes underscores to hyphens for OpenSSL and passes the string unchanged to GnuTLS (expecting the library to handle its own older variants) when processing lists of cipher suites in the **tls\_require\_ciphers** options (the global option and the *smtp* transport option).
- The **tls\_require\_ciphers** options operate differently, as described in the sections 42.4 and 42.5.
- The **tls\_dh\_min\_bits** SMTP transport option is only honoured by GnuTLS. When using OpenSSL, this option is ignored. (If an API is found to let OpenSSL be configured in this way, let the Exim Maintainers know and we'll likely use it).
- With GnuTLS, if an explicit list is used for the **tls\_privatekey** main option, it must be ordered to match the **tls\_certificate** list.
- Some other recently added features may only be available in one or the other. This should be documented with the feature. If the documentation does not explicitly state that the feature is infeasible in the other TLS implementation, then patches are welcome.

## 42.3 GnuTLS parameter computation

This section only applies if **tls\_dhparam** is set to `historic` or to an explicit path; if the latter, then the text about generation still applies, but not the chosen filename. By default, as of Exim 4.80 a hard-coded D-H prime is used. See the documentation of **tls\_dhparam** for more information.

GnuTLS uses D-H parameters that may take a substantial amount of time to compute. It is unreasonable to re-compute them for every TLS session. Therefore, Exim keeps this data in a file in its spool directory, called *gnutls-params-NNNN* for some value of NNNN, corresponding to the number of bits requested. The file is owned by the Exim user and is readable only by its owner. Every Exim process that start up GnuTLS reads the D-H parameters from this file. If the file does not exist, the first Exim process that needs it computes the data and writes it to a temporary file which is renamed once it is complete. It does not matter if several Exim processes do this simultaneously (apart from wasting a few resources). Once a file is in place, new Exim processes immediately start using it.

For maximum security, the parameters that are stored in this file should be recalculated periodically, the frequency depending on your paranoia level. If you are avoiding using the fixed D-H primes published in RFCs, then you are concerned about some advanced attacks and will wish to do this; if you do not regenerate then you might as well stick to the standard primes.

Arranging this is easy in principle; just delete the file when you want new values to be computed. However, there may be a problem. The calculation of new parameters needs random numbers, and these are obtained from */dev/random*. If the system is not very active, */dev/random* may delay returning data until enough randomness (entropy) is available. This may cause Exim to hang for a substantial amount of time, causing timeouts on incoming connections.

The solution is to generate the parameters externally to Exim. They are stored in *gnutls-params-N* in PEM format, which means that they can be generated externally using the *certtool* command that is part of GnuTLS.

To replace the parameters with new ones, instead of deleting the file and letting Exim re-create it, you can generate new parameters using *certtool* and, when this has been done, replace Exim's cache file by renaming. The relevant commands are something like this:

```
# ls
[ look for file; assume gnutls-params-2236 is the most recent ]
# rm -f new-params
# touch new-params
# chown exim:exim new-params
```

```
# chmod 0600 new-params
# certtool --generate-dh-params --bits 2236 >>new-params
# openssl dhparam -noout -text -in new-params | head
[ check the first line, make sure it's not more than 2236;
  if it is, then go back to the start ("rm") and repeat
  until the size generated is at most the size requested ]
# chmod 0400 new-params
# mv new-params gnutls-params-2236
```

If Exim never has to generate the parameters itself, the possibility of stalling is removed.

The filename changed in Exim 4.80, to gain the `-bits` suffix. The value which Exim will choose depends upon the version of GnuTLS in use. For older GnuTLS, the value remains hard-coded in Exim as 1024. As of GnuTLS 2.12.x, there is a way for Exim to ask for the "normal" number of bits for D-H public-key usage, and Exim does so. This attempt to remove Exim from TLS policy decisions failed, as GnuTLS 2.12 returns a value higher than the current hard-coded limit of the NSS library. Thus Exim gains the **`tls_dh_max_bits`** global option, which applies to all D-H usage, client or server. If the value returned by GnuTLS is greater than **`tls_dh_max_bits`** then the value will be clamped down to **`tls_dh_max_bits`**. The default value has been set at the current NSS limit, which is still much higher than Exim historically used.

The filename and bits used will change as the GnuTLS maintainers change the value for their parameter `GNUTLS_SEC_PARAM_NORMAL`, as clamped by **`tls_dh_max_bits`**. At the time of writing (mid 2012), GnuTLS 2.12 recommends 2432 bits, while NSS is limited to 2236 bits.

In fact, the requested value will be *\*lower\** than **`tls_dh_max_bits`**, to increase the chance of the generated prime actually being within acceptable bounds, as GnuTLS has been observed to overshoot. Note the check step in the procedure above. There is no sane procedure available to Exim to double-check the size of the generated prime, so it might still be too large.

## 42.4 Requiring specific ciphers in OpenSSL

There is a function in the OpenSSL library that can be passed a list of cipher suites before the cipher negotiation takes place. This specifies which ciphers are acceptable. The list is colon separated and may contain names like DES-CBC3-SHA. Exim passes the expanded value of **`tls_require_ciphers`** directly to this function call. Many systems will install the OpenSSL manual-pages, so you may have *ciphers(1)* available to you. The following quotation from the OpenSSL documentation specifies what forms of item are allowed in the cipher string:

- It can consist of a single cipher suite such as RC4-SHA.
- It can represent a list of cipher suites containing a certain algorithm, or cipher suites of a certain type. For example SHA1 represents all ciphers suites using the digest algorithm SHA1 and SSLv3 represents all SSL v3 algorithms.
- Lists of cipher suites can be combined in a single cipher string using the + character. This is used as a logical and operation. For example SHA1+DES represents all cipher suites containing the SHA1 and the DES algorithms.

Each cipher string can be optionally preceded by one of the characters `!`, `-` or `+`.

- If `!` is used, the ciphers are permanently deleted from the list. The ciphers deleted can never reappear in the list even if they are explicitly stated.
- If `-` is used, the ciphers are deleted from the list, but some or all of the ciphers can be added again by later options.
- If `+` is used, the ciphers are moved to the end of the list. This option does not add any new ciphers; it just moves matching existing ones.

If none of these characters is present, the string is interpreted as a list of ciphers to be appended to the current preference list. If the list includes any ciphers already present they will be ignored: that is, they will not be moved to the end of the list.

The OpenSSL *ciphers(1)* command may be used to test the results of a given string:

```
# note single-quotes to get ! past any shell history expansion
$ openssl ciphers 'HIGH:!MD5:!SHA1'
```

This example will let the library defaults be permitted on the MX port, where there's probably no identity verification anyway, but ups the ante on the submission ports where the administrator might have some influence on the choice of clients used:

```
# OpenSSL variant; see man ciphers(1)
tls_require_ciphers = ${if =={$received_port}{25}\
    {DEFAULT}\
    {HIGH:!MD5:!SHA1}}
```

This example will prefer ECDSA-authenticated ciphers over RSA ones:

```
tls_require_ciphers = ECDSA:RSA:!COMPLEMENTOFDEFAULT
```

## 42.5 Requiring specific ciphers or other parameters in GnuTLS

The GnuTLS library allows the caller to provide a "priority string", documented as part of the *gnutls\_priority\_init* function. This is very similar to the ciphersuite specification in OpenSSL.

The **tls\_require\_ciphers** option is treated as the GnuTLS priority string and controls both protocols and ciphers.

The **tls\_require\_ciphers** option is available both as an global option, controlling how Exim behaves as a server, and also as an option of the *smtp* transport, controlling how Exim behaves as a client. In both cases the value is string expanded. The resulting string is not an Exim list and the string is given to the GnuTLS library, so that Exim does not need to be aware of future feature enhancements of GnuTLS.

Documentation of the strings accepted may be found in the GnuTLS manual, under "Priority strings". This is online as [http://www.gnutls.org/manual/html\\_node/Priority-Strings.html](http://www.gnutls.org/manual/html_node/Priority-Strings.html), but beware that this relates to GnuTLS 3, which may be newer than the version installed on your system. If you are using GnuTLS 3, then the example code <http://www.gnutls.org/manual/gnutls.html#Listing-the-ciphersuites-in-a-priority-string> on that site can be used to test a given string.

For example:

```
# Disable older versions of protocols
tls_require_ciphers = NORMAL:%LATEST_RECORD_VERSION:-VERS-SSL3.0
```

Prior to Exim 4.80, an older API of GnuTLS was used, and Exim supported three additional options, "**gnutls\_require\_kx**", "**gnutls\_require\_mac**" and "**gnutls\_require\_protocols**". **tls\_require\_ciphers** was an Exim list.

This example will let the library defaults be permitted on the MX port, where there's probably no identity verification anyway, and lowers security further by increasing compatibility; but this ups the ante on the submission ports where the administrator might have some influence on the choice of clients used:

```
# GnuTLS variant
tls_require_ciphers = ${if =={$received_port}{25}\
    {NORMAL:%COMPAT}\
    {SECURE128}}
```

## 42.6 Configuring an Exim server to use TLS

When Exim has been built with TLS support, it advertises the availability of the STARTTLS command to client hosts that match **tls\_advertise\_hosts**, but not to any others. The default value of this option is \*, which means that STARTTLS is always advertised. Set it to blank to never advertise; this is reasonable for systems that want to use TLS only as a client.

If STARTTLS is to be used you need to set some other options in order to make TLS available.

If a client issues a STARTTLS command and there is some configuration problem in the server, the command is rejected with a 454 error. If the client persists in trying to issue SMTP commands, all except QUIT are rejected with the error

554 Security failure

If a STARTTLS command is issued within an existing TLS session, it is rejected with a 554 error code.

To enable TLS operations on a server, the **tls\_advertise\_hosts** option must be set to match some hosts. The default is `*` which matches all hosts.

If this is all you do, TLS encryption will be enabled but not authentication - meaning that the peer has no assurance it is actually you he is talking to. You gain protection from a passive sniffer listening on the wire but not from someone able to intercept the communication.

Further protection requires some further configuration at the server end.

To make TLS work you need to set, in the server,

```
tls_certificate = /some/file/name
tls_privatekey = /some/file/name
```

These options are, in fact, expanded strings, so you can make them depend on the identity of the client that is connected if you wish. The first file contains the server's X509 certificate, and the second contains the private key that goes with it. These files need to be PEM format and readable by the Exim user, and must always be given as full path names. The key must not be password-protected. They can be the same file if both the certificate and the key are contained within it. If **tls\_privatekey** is not set, or if its expansion is forced to fail or results in an empty string, this is assumed to be the case. The certificate file may also contain intermediate certificates that need to be sent to the client to enable it to authenticate the server's certificate.

For dual-stack (eg. RSA and ECDSA) configurations, these options can be colon-separated lists of file paths. Ciphers using given authentication algorithms require the presence of a suitable certificate to supply the public-key. The server selects among the certificates to present to the client depending on the selected cipher, hence the priority ordering for ciphers will affect which certificate is used.

If you do not understand about certificates and keys, please try to find a source of this background information, which is not Exim-specific. (There are a few comments below in section 42.12.)

**Note:** These options do not apply when Exim is operating as a client – they apply only in the case of a server. If you need to use a certificate in an Exim client, you must set the options of the same names in an *smtp* transport.

With just these options, an Exim server will be able to use TLS. It does not require the client to have a certificate (but see below for how to insist on this). There is one other option that may be needed in other situations. If

```
tls_dhparam = /some/file/name
```

is set, the SSL library is initialized for the use of Diffie-Hellman ciphers with the parameters contained in the file. Set this to `none` to disable use of DH entirely, by making no prime available:

```
tls_dhparam = none
```

This may also be set to a string identifying a standard prime to be used for DH; if it is set to `default` or, for OpenSSL, is unset, then the prime used is `ike23`. There are a few standard primes available, see the documentation for **tls\_dhparam** for the complete list.

See the command

```
openssl dhparam
```

for a way of generating file data.

The strings supplied for these three options are expanded every time a client host connects. It is therefore possible to use different certificates and keys for different hosts, if you so wish, by making

use of the client's IP address in *\$sender\_host\_address* to control the expansion. If a string expansion is forced to fail, Exim behaves as if the option is not set.

The variable *\$tls\_in\_cipher* is set to the cipher suite that was negotiated for an incoming TLS connection. It is included in the *Received:* header of an incoming message (by default – you can, of course, change this), and it is also included in the log line that records a message's arrival, keyed by “X=”, unless the **tls\_cipher** log selector is turned off. The **encrypted** condition can be used to test for specific cipher suites in ACLs.

Once TLS has been established, the ACLs that run for subsequent SMTP commands can check the name of the cipher suite and vary their actions accordingly. The cipher suite names vary, depending on which TLS library is being used. For example, OpenSSL uses the name DES-CBC3-SHA for the cipher suite which in other contexts is known as TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA. Check the OpenSSL or GnuTLS documentation for more details.

For outgoing SMTP deliveries, *\$tls\_out\_cipher* is used and logged (again depending on the **tls\_cipher** log selector).

## 42.7 Requesting and verifying client certificates

If you want an Exim server to request a certificate when negotiating a TLS session with a client, you must set either **tls\_verify\_hosts** or **tls\_try\_verify\_hosts**. You can, of course, set either of them to \* to apply to all TLS connections. For any host that matches one of these options, Exim requests a certificate as part of the setup of the TLS session. The contents of the certificate are verified by comparing it with a list of expected certificates. These may be the system default set (depending on library version), an explicit file or, depending on library version, a directory, identified by **tls\_verify\_certificates**.

A file can contain multiple certificates, concatenated end to end. If a directory is used (OpenSSL only), each certificate must be in a separate file, with a name (or a symbolic link) of the form *<hash>.0*, where *<hash>* is a hash value constructed from the certificate. You can compute the relevant hash by running the command

```
openssl x509 -hash -noout -in /cert/file
```

where */cert/file* contains a single certificate.

The difference between **tls\_verify\_hosts** and **tls\_try\_verify\_hosts** is what happens if the client does not supply a certificate, or if the certificate does not match any of the certificates in the collection named by **tls\_verify\_certificates**. If the client matches **tls\_verify\_hosts**, the attempt to set up a TLS session is aborted, and the incoming connection is dropped. If the client matches **tls\_try\_verify\_hosts**, the (encrypted) SMTP session continues. ACLs that run for subsequent SMTP commands can detect the fact that no certificate was verified, and vary their actions accordingly. For example, you can insist on a certificate before accepting a message for relaying, but not when the message is destined for local delivery.

When a client supplies a certificate (whether it verifies or not), the value of the Distinguished Name of the certificate is made available in the variable *\$tls\_in\_peerdn* during subsequent processing of the message.

Because it is often a long text string, it is not included in the log line or *Received:* header by default. You can arrange for it to be logged, keyed by “DN=”, by setting the **tls\_peerdn** log selector, and you can use **received\_header\_text** to change the *Received:* header. When no certificate is supplied, *\$tls\_in\_peerdn* is empty.

## 42.8 Revoked certificates

Certificate issuing authorities issue Certificate Revocation Lists (CRLs) when certificates are revoked. If you have such a list, you can pass it to an Exim server using the global option called **tls\_crl** and to an Exim client using an identically named option for the *smtp* transport. In each case, the value of the option is expanded and must then be the name of a file that contains a CRL in PEM format. The

downside is that clients have to periodically re-download a potentially huge file from every certificate authority they know of.

The way with most moving parts at query time is Online Certificate Status Protocol (OCSP), where the client verifies the certificate against an OCSP server run by the CA. This lets the CA track all usage of the certs. It requires running software with access to the private key of the CA, to sign the responses to the OCSP queries. OCSP is based on HTTP and can be proxied accordingly.

The only widespread OCSP server implementation (known to this writer) comes as part of OpenSSL and aborts on an invalid request, such as connecting to the port and then disconnecting. This requires re-entering the passphrase each time some random client does this.

The third way is OCSP Stapling; in this, the server using a certificate issued by the CA periodically requests an OCSP proof of validity from the OCSP server, then serves it up inline as part of the TLS negotiation. This approach adds no extra round trips, does not let the CA track users, scales well with number of certs issued by the CA and is resilient to temporary OCSP server failures, as long as the server starts retrying to fetch an OCSP proof some time before its current proof expires. The downside is that it requires server support.

Unless Exim is built with the support disabled, or with GnuTLS earlier than version 3.3.16 / 3.4.8 support for OCSP stapling is included.

There is a global option called **tls\_ocsp\_file**. The file specified therein is expected to be in DER format, and contain an OCSP proof. Exim will serve it as part of the TLS handshake. This option will be re-expanded for SNI, if the **tls\_certificate** option contains **tls\_in\_sni**, as per other TLS options.

Exim does not at this time implement any support for fetching a new OCSP proof. The burden is on the administrator to handle this, outside of Exim. The file specified should be replaced atomically, so that the contents are always valid. Exim will expand the **tls\_ocsp\_file** option on each connection, so a new file will be handled transparently on the next connection.

When built with OpenSSL Exim will check for a valid next update timestamp in the OCSP proof; if not present, or if the proof has expired, it will be ignored.

For the client to be able to verify the stapled OCSP the server must also supply, in its stapled information, any intermediate certificates for the chain leading to the OCSP proof from the signer of the server certificate. There may be zero or one such. These intermediate certificates should be added to the server OCSP stapling file named by **tls\_ocsp\_file**.

Note that the proof only covers the terminal server certificate, not any of the chain from CA to it.

There is no current way to staple a proof for a client certificate.

```
A helper script "ocsp_fetch.pl" for fetching a proof from a CA
OCSP server is supplied. The server URL may be included in the
server certificate, if the CA is helpful.
```

```
One failure mode seen was the OCSP Signer cert expiring before the end
of validity of the OCSP proof. The checking done by Exim/OpenSSL
noted this as invalid overall, but the re-fetch script did not.
```

## 42.9 Configuring an Exim client to use TLS

The **tls\_cipher** and **tls\_peerdn** log selectors apply to outgoing SMTP deliveries as well as to incoming, the latter one causing logging of the server certificate's DN. The remaining client configuration for TLS is all within the *smtp* transport.

It is not necessary to set any options to have TLS work in the *smtp* transport. If Exim is built with TLS support, and TLS is advertised by a server, the *smtp* transport always tries to start a TLS session. However, this can be prevented by setting **hosts\_avoid\_tls** (an option of the transport) to a list of server hosts for which TLS should not be used.

If you do not want Exim to attempt to send messages unencrypted when an attempt to set up an encrypted connection fails in any way, you can set **hosts\_require\_tls** to a list of hosts for which encryption is mandatory. For those hosts, delivery is always deferred if an encrypted connection cannot be set up. If there are any other hosts for the address, they are tried in the usual way.

When the server host is not in **hosts\_require\_tls**, Exim may try to deliver the message unencrypted. It always does this if the response to STARTTLS is a 5xx code. For a temporary error code, or for a failure to negotiate a TLS session after a success response code, what happens is controlled by the **tls\_tempfail\_tryclear** option of the *smtp* transport. If it is false, delivery to this host is deferred, and other hosts (if available) are tried. If it is true, Exim attempts to deliver unencrypted after a 4xx response to STARTTLS, and if STARTTLS is accepted, but the subsequent TLS negotiation fails, Exim closes the current connection (because it is in an unknown state), opens a new one to the same host, and then tries the delivery unencrypted.

The **tls\_certificate** and **tls\_privatekey** options of the *smtp* transport provide the client with a certificate, which is passed to the server if it requests it. If the server is Exim, it will request a certificate only if **tls\_verify\_hosts** or **tls\_try\_verify\_hosts** matches the client.

If the **tls\_verify\_certificates** option is set on the *smtp* transport, it specifies a collection of expected server certificates. These may be the system default set (depending on library version), a file, or (depending on library version) a directory. The client verifies the server's certificate against this collection, taking into account any revoked certificates that are in the list defined by **tls\_crl**. Failure to verify fails the TLS connection unless either of the **tls\_verify\_hosts** or **tls\_try\_verify\_hosts** options are set.

The **tls\_verify\_hosts** and **tls\_try\_verify\_hosts** options restrict certificate verification to the listed servers. Verification either must or need not succeed respectively.

The *smtp* transport has two OCSP-related options: **hosts\_require\_ocsp**; a host-list for which a Certificate Status is requested and required for the connection to proceed. The default value is empty. **hosts\_request\_ocsp**; a host-list for which (additionally) a Certificate Status is requested (but not necessarily verified). The default value is "\*" meaning that requests are made unless configured otherwise.

The host(s) should also be in **hosts\_require\_tls**, and **tls\_verify\_certificates** configured for the transport, for OCSP to be relevant.

If **tls\_require\_ciphers** is set on the *smtp* transport, it must contain a list of permitted cipher suites. If either of these checks fails, delivery to the current host is abandoned, and the *smtp* transport tries to deliver to alternative hosts, if any.

**Note:** These options must be set in the *smtp* transport for Exim to use TLS when it is operating as a client. Exim does not assume that a server certificate (set by the global options of the same name) should also be used when operating as a client.

All the TLS options in the *smtp* transport are expanded before use, with *\$host* and *\$host\_address* containing the name and address of the server to which the client is connected. Forced failure of an expansion causes Exim to behave as if the relevant option were unset.

Before an SMTP connection is established, the *\$tls\_out\_bits*, *\$tls\_out\_cipher*, *\$tls\_out\_peerdn* and *\$tls\_out\_sni* variables are emptied. (Until the first connection, they contain the values that were set when the message was received.) If STARTTLS is subsequently successfully obeyed, these variables are set to the relevant values for the outgoing connection.

## 42.10 Use of TLS Server Name Indication

With TLS1.0 or above, there is an extension mechanism by which extra information can be included at various points in the protocol. One of these extensions, documented in RFC 6066 (and before that RFC 4366) is "Server Name Indication", commonly "SNI". This extension is sent by the client in the initial handshake, so that the server can examine the servername within and possibly choose to use different certificates and keys (and more) for this session.

This is analogous to HTTP's "Host:" header, and is the main mechanism by which HTTPS-enabled web-sites can be virtual-hosted, many sites to one IP address.

With SMTP to MX, there are the same problems here as in choosing the identity against which to validate a certificate: you can't rely on insecure DNS to provide the identity which you then cryptographically verify. So this will be of limited use in that environment.

With SMTP to Submission, there is a well-defined hostname which clients are connecting to and can validate certificates against. Thus clients **can** choose to include this information in the TLS negotiation. If this becomes wide-spread, then hosters can choose to present different certificates to different clients. Or even negotiate different cipher suites.

The **tls\_sni** option on an SMTP transport is an expanded string; the result, if not empty, will be sent on a TLS session as part of the handshake. There's nothing more to it. Choosing a sensible value not derived insecurely is the only point of caution. The *\$tls\_out\_sni* variable will be set to this string for the lifetime of the client connection (including during authentication).

Except during SMTP client sessions, if *\$tls\_in\_sni* is set then it is a string received from a client. It can be logged with the **log\_selector** item *+tls\_sni*.

If the string *tls\_in\_sni* appears in the main section's **tls\_certificate** option (prior to expansion) then the following options will be re-expanded during TLS session handshake, to permit alternative values to be chosen:

- **tls\_certificate**
- **tls\_crl**
- **tls\_privatekey**
- **tls\_verify\_certificates**
- **tls\_ocsp\_file**

Great care should be taken to deal with matters of case, various injection attacks in the string (*. /* or *SQL*), and ensuring that a valid filename can always be referenced; it is important to remember that *\$tls\_in\_sni* is arbitrary unverified data provided prior to authentication. Further, the initial certificate is loaded before SNI is arrived, so an expansion for **tls\_certificate** must have a default which is used when *\$tls\_in\_sni* is empty.

The Exim developers are proceeding cautiously and so far no other TLS options are re-expanded.

When Exim is built against OpenSSL, OpenSSL must have been built with support for TLS Extensions. This holds true for OpenSSL 1.0.0+ and 0.9.8+ with *enable-tlsex* in *EXTRACONFIGURE*. If you invoke *openssl s\_client -h* and see *-servername* in the output, then OpenSSL has support.

When Exim is built against GnuTLS, SNI support is available as of GnuTLS 0.5.10. (Its presence predates the current API which Exim uses, so if Exim built, then you have SNI support).

## 42.11 Multiple messages on the same encrypted TCP/IP connection

Exim sends multiple messages down the same TCP/IP connection by starting up an entirely new delivery process for each message, passing the socket from one process to the next. This implementation does not fit well with the use of TLS, because there is quite a lot of state information associated with a TLS connection, not just a socket identification. Passing all the state information to a new process is not feasible. Consequently, for sending using TLS Exim starts an additional proxy process for handling the encryption, piping the unencrypted data stream from and to the delivery processes.

An older mode of operation can be enabled on a per-host basis by the **hosts\_noproxy\_tls** option on the *smtp* transport. If the host matches this list the proxy process described above is not used; instead Exim

shuts down an existing TLS session being run by the delivery process before passing the socket to a new process. The new process may then try to start a new TLS session, and if successful, may try to re-authenticate if AUTH is in use, before sending the next message.



The RFC is not clear as to whether or not an SMTP session continues in clear after TLS has been shut down, or whether TLS may be restarted again later, as just described. However, if the server is Exim, this shutdown and reinitialization works. It is not known which (if any) other servers operate successfully if the client closes a TLS session and continues with unencrypted SMTP, but there are certainly some that do not work. For such servers, Exim should not pass the socket to another process, because the failure of the subsequent attempt to use it would cause Exim to record a temporary host error, and delay other deliveries to that host.

To test for this case, Exim sends an EHLO command to the server after closing down the TLS session. If this fails in any way, the connection is closed instead of being passed to a new delivery process, but no retry information is recorded.

There is also a manual override; you can set **hosts\_nopass\_tls** on the *smtp* transport to match those hosts for which Exim should not pass connections to new processes if TLS has been used.

## 42.12 Certificates and all that

In order to understand fully how TLS works, you need to know about certificates, certificate signing, and certificate authorities. This is not the place to give a tutorial, especially as I do not know very much about it myself. Some helpful introduction can be found in the FAQ for the SSL addition to Apache, currently at

[http://www.modssl.org/docs/2.7/ssl\\_faq.html#ToC24](http://www.modssl.org/docs/2.7/ssl_faq.html#ToC24)

Other parts of the *modssl* documentation are also helpful, and have links to further files. Eric Rescorla's book, *SSL and TLS*, published by Addison-Wesley (ISBN 0-201-61598-3), contains both introductory and more in-depth descriptions. Some sample programs taken from the book are available from

<http://www.rtfm.com/openssl-examples/>

## 42.13 Certificate chains

The file named by **tls\_certificate** may contain more than one certificate. This is useful in the case where the certificate that is being sent is validated by an intermediate certificate which the other end does not have. Multiple certificates must be in the correct order in the file. First the host's certificate itself, then the first intermediate certificate to validate the issuer of the host certificate, then the next intermediate certificate to validate the issuer of the first intermediate certificate, and so on, until finally (optionally) the root certificate. The root certificate must already be trusted by the recipient for validation to succeed, of course, but if it's not preinstalled, sending the root certificate along with the rest makes it available for the user to install if the receiving end is a client MUA that can interact with a user.

Note that certificates using MD5 are unlikely to work on today's Internet; even if your libraries allow loading them for use in Exim when acting as a server, increasingly clients will not accept such certificates. The error diagnostics in such a case can be frustratingly vague.

## 42.14 Self-signed certificates

You can create a self-signed certificate using the *req* command provided with OpenSSL, like this:

```
openssl req -x509 -newkey rsa:1024 -keyout file1 -out file2 \  
-days 9999 -nodes
```

*file1* and *file2* can be the same file; the key and the certificate are delimited and so can be identified independently. The **-days** option specifies a period for which the certificate is valid. The **-nodes** option is important: if you do not set it, the key is encrypted with a passphrase that you are prompted for, and any use that is made of the key causes more prompting for the passphrase. This is not helpful if you are going to use this certificate and key in an MTA, where prompting is not possible.

NB: we are now past the point where 9999 days takes us past the 32-bit Unix epoch. If your system uses unsigned *time\_t* (most do) and is 32-bit, then the above command might produce a date in the

past. Think carefully about the lifetime of the systems you're deploying, and either reduce the duration of the certificate or reconsider your platform deployment. (At time of writing, reducing the duration is the most likely choice, but the inexorable progression of time takes us steadily towards an era where this will not be a sensible resolution).

A self-signed certificate made in this way is sufficient for testing, and may be adequate for all your requirements if you are mainly interested in encrypting transfers, and not in secure identification.

However, many clients require that the certificate presented by the server be a user (also called "leaf" or "site") certificate, and not a self-signed certificate. In this situation, the self-signed certificate described above must be installed on the client host as a trusted root *certification authority* (CA), and the certificate used by Exim must be a user certificate signed with that self-signed certificate.

For information on creating self-signed CA certificates and using them to sign user certificates, see the *General implementation overview* chapter of the Open-source PKI book, available online at <http://ospkibook.sourceforge.net/>.

## 43. Access control lists

Access Control Lists (ACLs) are defined in a separate section of the run time configuration file, headed by “begin acl”. Each ACL definition starts with a name, terminated by a colon. Here is a complete ACL section that contains just one very small ACL:

```
begin acl
small_acl:
accept    hosts = one.host.only
```

You can have as many lists as you like in the ACL section, and the order in which they appear does not matter. The lists are self-terminating.

The majority of ACLs are used to control Exim's behaviour when it receives certain SMTP commands. This applies both to incoming TCP/IP connections, and when a local process submits a message using SMTP by specifying the **-bs** option. The most common use is for controlling which recipients are accepted in incoming messages. In addition, you can define an ACL that is used to check local non-SMTP messages. The default configuration file contains an example of a realistic ACL for checking RCPT commands. This is discussed in chapter 7.

### 43.1 Testing ACLs

The **-bh** command line option provides a way of testing your ACL configuration locally by running a fake SMTP session with which you interact.

### 43.2 Specifying when ACLs are used

In order to cause an ACL to be used, you have to name it in one of the relevant options in the main part of the configuration. These options are:

<b>acl_not_smtp</b>	ACL for non-SMTP messages
<b>acl_not_smtp_mime</b>	ACL for non-SMTP MIME parts
<b>acl_not_smtp_start</b>	ACL at start of non-SMTP message
<b>acl_smtp_auth</b>	ACL for AUTH
<b>acl_smtp_connect</b>	ACL for start of SMTP connection
<b>acl_smtp_data</b>	ACL after DATA is complete
<b>acl_smtp_data_prdr</b>	ACL for each recipient, after DATA is complete
<b>acl_smtp_dkim</b>	ACL for each DKIM signer
<b>acl_smtp_etrn</b>	ACL for ETRN
<b>acl_smtp_expn</b>	ACL for EXPN
<b>acl_smtp_helo</b>	ACL for HELO or EHLO
<b>acl_smtp_mail</b>	ACL for MAIL
<b>acl_smtp_mailauth</b>	ACL for the AUTH parameter of MAIL
<b>acl_smtp_mime</b>	ACL for content-scanning MIME parts
<b>acl_smtp_notquit</b>	ACL for non-QUIT terminations
<b>acl_smtp_predata</b>	ACL at start of DATA command
<b>acl_smtp_quit</b>	ACL for QUIT
<b>acl_smtp_rcpt</b>	ACL for RCPT
<b>acl_smtp_starttls</b>	ACL for STARTTLS
<b>acl_smtp_vrfy</b>	ACL for VRFY

For example, if you set

```
acl_smtp_rcpt = small_acl
```

the little ACL defined above is used whenever Exim receives a RCPT command in an SMTP dialogue. The majority of policy tests on incoming messages can be done when RCPT commands arrive. A rejection of RCPT should cause the sending MTA to give up on the recipient address contained in the RCPT command, whereas rejection at other times may cause the client MTA to keep on trying to deliver the message. It is therefore recommended that you do as much testing as possible at RCPT time.

### 43.3 The non-SMTP ACLs

The non-SMTP ACLs apply to all non-interactive incoming messages, that is, they apply to batched SMTP as well as to non-SMTP messages. (Batched SMTP is not really SMTP.) Many of the ACL conditions (for example, host tests, and tests on the state of the SMTP connection such as encryption and authentication) are not relevant and are forbidden in these ACLs. However, the sender and recipients are known, so the **senders** and **sender\_domains** conditions and the *\$sender\_address* and *\$recipients* variables can be used. Variables such as *\$authenticated\_sender* are also available. You can specify added header lines in any of these ACLs.

The **acl\_not\_smtp\_start** ACL is run right at the start of receiving a non-SMTP message, before any of the message has been read. (This is the analogue of the **acl\_smtp\_predata** ACL for SMTP input.) In the case of batched SMTP input, it runs after the DATA command has been reached. The result of this ACL is ignored; it cannot be used to reject a message. If you really need to, you could set a value in an ACL variable here and reject based on that in the **acl\_not\_smtp** ACL. However, this ACL can be used to set controls, and in particular, it can be used to set

```
control = suppress_local_fixups
```

This cannot be used in the other non-SMTP ACLs because by the time they are run, it is too late.

The **acl\_not\_smtp\_mime** ACL is available only when Exim is compiled with the content-scanning extension. For details, see chapter 44.

The **acl\_not\_smtp** ACL is run just before the *local\_scan()* function. Any kind of rejection is treated as permanent, because there is no way of sending a temporary error for these kinds of message.

### 43.4 The SMTP connect ACL

The ACL test specified by **acl\_smtp\_connect** happens at the start of an SMTP session, after the test specified by **host\_reject\_connection** (which is now an anomaly) and any TCP Wrappers testing (if configured). If the connection is accepted by an **accept** verb that has a **message** modifier, the contents of the message override the banner message that is otherwise specified by the **smtp\_banner** option.

### 43.5 The EHLO/HELO ACL

The ACL test specified by **acl\_smtp\_helo** happens when the client issues an EHLO or HELO command, after the tests specified by **helo\_accept\_junk\_hosts**, **helo\_allow\_chars**, **helo\_verify\_hosts**, and **helo\_try\_verify\_hosts**. Note that a client may issue more than one EHLO or HELO command in an SMTP session, and indeed is required to issue a new EHLO or HELO after successfully setting up encryption following a STARTTLS command.

Note also that a deny neither forces the client to go away nor means that mail will be refused on the connection. Consider checking for *\$sender\_helo\_name* being defined in a MAIL or RCPT ACL to do that.

If the command is accepted by an **accept** verb that has a **message** modifier, the message may not contain more than one line (it will be truncated at the first newline and a panic logged if it does). Such a message cannot affect the EHLO options that are listed on the second and subsequent lines of an EHLO response.

### 43.6 The DATA ACLs

Two ACLs are associated with the DATA command, because it is two-stage command, with two responses being sent to the client. When the DATA command is received, the ACL defined by **acl\_smtp\_predata** is obeyed. This gives you control after all the RCPT commands, but before the message itself is received. It offers the opportunity to give a negative response to the DATA command before the data is transmitted. Header lines added by MAIL or RCPT ACLs are not visible at this time, but any that are defined here are visible when the **acl\_smtp\_data** ACL is run.

You cannot test the contents of the message, for example, to verify addresses in the headers, at RCPT time or when the DATA command is received. Such tests have to appear in the ACL that is run after

the message itself has been received, before the final response to the DATA command is sent. This is the ACL specified by **acl\_smtp\_data**, which is the second ACL that is associated with the DATA command.

If CHUNKING was advertised and a BDAT command sequence is received, the **acl\_smtp\_predata** ACL is not run. The **acl\_smtp\_data** is run after the last BDAT command and all of the data specified is received.

For both of these ACLs, it is not possible to reject individual recipients. An error response rejects the entire message. Unfortunately, it is known that some MTAs do not treat hard (5xx) responses to the DATA command (either before or after the data) correctly – they keep the message on their queues and try again later, but that is their problem, though it does waste some of your resources.

The **acl\_smtp\_data** ACL is run after the **acl\_smtp\_data\_prdr**, the **acl\_smtp\_dkim** and the **acl\_smtp\_mime** ACLs.

### 43.7 The SMTP DKIM ACL

The **acl\_smtp\_dkim** ACL is available only when Exim is compiled with DKIM support enabled (which is the default).

The ACL test specified by **acl\_smtp\_dkim** happens after a message has been received, and is executed for each DKIM signature found in a message. If not otherwise specified, the default action is to accept.

This ACL is evaluated before **acl\_smtp\_mime** and **acl\_smtp\_data**.

For details on the operation of DKIM, see chapter 57.

### 43.8 The SMTP MIME ACL

The **acl\_smtp\_mime** option is available only when Exim is compiled with the content-scanning extension. For details, see chapter 44.

This ACL is evaluated after **acl\_smtp\_dkim** but before **acl\_smtp\_data**.

### 43.9 The SMTP PRDR ACL

The **acl\_smtp\_data\_prdr** ACL is available only when Exim is compiled with PRDR support enabled (which is the default). It becomes active only when the PRDR feature is negotiated between client and server for a message, and more than one recipient has been accepted.

The ACL test specified by **acl\_smtp\_data\_prdr** happens after a message has been received, and is executed once for each recipient of the message with *\$local\_part* and *\$domain* valid. The test may accept, defer or deny for individual recipients. The **acl\_smtp\_data** will still be called after this ACL and can reject the message overall, even if this ACL has accepted it for some or all recipients.

PRDR may be used to support per-user content filtering. Without it one must defer any recipient after the first that has a different content-filter configuration. With PRDR, the RCPT-time check for this can be disabled when the variable *\$prdr\_requested* is “yes”. Any required difference in behaviour of the main DATA-time ACL should however depend on the PRDR-time ACL having run, as Exim will avoid doing so in some situations (e.g. single-recipient mails).

See also the **prdr\_enable** global option and the **hosts\_try\_prdr** smtp transport option.

This ACL is evaluated after **acl\_smtp\_dkim** but before **acl\_smtp\_data**. If the ACL is not defined, processing completes as if the feature was not requested by the client.

### 43.10 The QUIT ACL

The ACL for the SMTP QUIT command is anomalous, in that the outcome of the ACL does not affect the response code to QUIT, which is always 221. Thus, the ACL does not in fact control any access. For this reason, it may only accept or warn as its final result.

This ACL can be used for tasks such as custom logging at the end of an SMTP session. For example, you can use ACL variables in other ACLs to count messages, recipients, etc., and log the totals at QUIT time using one or more **logwrite** modifiers on a **warn** verb.

**Warning:** Only the *\$acl\_cx* variables can be used for this, because the *\$acl\_mx* variables are reset at the end of each incoming message.

You do not need to have a final **accept**, but if you do, you can use a **message** modifier to specify custom text that is sent as part of the 221 response to QUIT.

This ACL is run only for a “normal” QUIT. For certain kinds of disastrous failure (for example, failure to open a log file, or when Exim is bombing out because it has detected an unrecoverable error), all SMTP commands from the client are given temporary error responses until QUIT is received or the connection is closed. In these special cases, the QUIT ACL does not run.

## 43.11 The not-QUIT ACL

The not-QUIT ACL, specified by **acl\_smtp\_notquit**, is run in most cases when an SMTP session ends without sending QUIT. However, when Exim itself is in bad trouble, such as being unable to write to its log files, this ACL is not run, because it might try to do things (such as write to log files) that make the situation even worse.

Like the QUIT ACL, this ACL is provided to make it possible to do customized logging or to gather statistics, and its outcome is ignored. The **delay** modifier is forbidden in this ACL, and the only permitted verbs are **accept** and **warn**.

When the not-QUIT ACL is running, the variable *\$smtp\_notquit\_reason* is set to a string that indicates the reason for the termination of the SMTP connection. The possible values are:

acl-drop	Another ACL issued a <b>drop</b> command
bad-commands	Too many unknown or non-mail commands
command-timeout	Timeout while reading SMTP commands
connection-lost	The SMTP connection has been lost
data-timeout	Timeout while reading message data
local-scan-error	The <i>local_scan()</i> function crashed
local-scan-timeout	The <i>local_scan()</i> function timed out
signal-exit	SIGTERM or SIGINT
synchronization-error	SMTP synchronization error
tls-failed	TLS failed to start

In most cases when an SMTP connection is closed without having received QUIT, Exim sends an SMTP response message before actually closing the connection. With the exception of the **acl-drop** case, the default message can be overridden by the **message** modifier in the not-QUIT ACL. In the case of a **drop** verb in another ACL, it is the message from the other ACL that is used.

## 43.12 Finding an ACL to use

The value of an **acl\_smtp\_XXX** option is expanded before use, so you can use different ACLs in different circumstances. For example,

```
acl_smtp_rcpt = ${if =25}${interface_port} \
                {acl_check_rcpt} {acl_check_rcpt_submit} }
```

In the default configuration file there are some example settings for providing an RFC 4409 message submission service on port 587 and a non-standard “smtps” service on port 465. You can use a string expansion like this to choose an ACL for MUAs on these ports which is more appropriate for this purpose than the default ACL on port 25.

The expanded string does not have to be the name of an ACL in the configuration file; there are other possibilities. Having expanded the string, Exim searches for an ACL as follows:

- If the string begins with a slash, Exim uses it as a file name, and reads its contents as an ACL. The lines are processed in the same way as lines in the Exim configuration file. In particular, continu-

ation lines are supported, blank lines are ignored, as are lines whose first non-whitespace character is “#”. If the file does not exist or cannot be read, an error occurs (typically causing a temporary failure of whatever caused the ACL to be run). For example:

```
acl_smtp_data = /etc/acls/\
    ${lookup{$sender_host_address}lsearch\
    /etc/acllist}{$value}{default}}
```

This looks up an ACL file to use on the basis of the host’s IP address, falling back to a default if the lookup fails. If an ACL is successfully read from a file, it is retained in memory for the duration of the Exim process, so that it can be re-used without having to re-read the file.

- If the string does not start with a slash, and does not contain any spaces, Exim searches the ACL section of the configuration for an ACL whose name matches the string.
- If no named ACL is found, or if the string contains spaces, Exim parses the string as an inline ACL. This can save typing in cases where you just want to have something like

```
acl_smtp_vrfy = accept
```

in order to allow free use of the VRFY command. Such a string may contain newlines; it is processed in the same way as an ACL that is read from a file.

### 43.13 ACL return codes

Except for the QUIT ACL, which does not affect the SMTP return code (see section 43.10 above), the result of running an ACL is either “accept” or “deny”, or, if some test cannot be completed (for example, if a database is down), “defer”. These results cause 2xx, 5xx, and 4xx return codes, respectively, to be used in the SMTP dialogue. A fourth return, “error”, occurs when there is an error such as invalid syntax in the ACL. This also causes a 4xx return code.

For the non-SMTP ACL, “defer” and “error” are treated in the same way as “deny”, because there is no mechanism for passing temporary errors to the submitters of non-SMTP messages.

ACLs that are relevant to message reception may also return “discard”. This has the effect of “accept”, but causes either the entire message or an individual recipient address to be discarded. In other words, it is a blackholing facility. Use it with care.

If the ACL for MAIL returns “discard”, all recipients are discarded, and no ACL is run for subsequent RCPT commands. The effect of “discard” in a RCPT ACL is to discard just the one recipient address. If there are no recipients left when the message’s data is received, the DATA ACL is not run. A “discard” return from the DATA or the non-SMTP ACL discards all the remaining recipients. The “discard” return is not permitted for the **acl\_smtp\_predata** ACL.

If the ACL for VRFY returns “accept”, a recipient verify (without callout) is done on the address and the result determines the SMTP response.

The *local\_scan()* function is always run, even if there are no remaining recipients; it may create new recipients.

### 43.14 Unset ACL options

The default actions when any of the **acl\_**xxx options are unset are not all the same. **Note:** These defaults apply only when the relevant ACL is not defined at all. For any defined ACL, the default action when control reaches the end of the ACL statements is “deny”.

For **acl\_smtp\_quit** and **acl\_not\_smtp\_start** there is no default because these two are ACLs that are used only for their side effects. They cannot be used to accept or reject anything.

For **acl\_not\_smtp**, **acl\_smtp\_auth**, **acl\_smtp\_connect**, **acl\_smtp\_data**, **acl\_smtp\_helo**, **acl\_smtp\_mail**, **acl\_smtp\_mailauth**, **acl\_smtp\_mime**, **acl\_smtp\_predata**, and **acl\_smtp\_starttls**, the action when the ACL is not defined is “accept”.

For the others (**acl\_smtp\_etrn**, **acl\_smtp\_expn**, **acl\_smtp\_rcpt**, and **acl\_smtp\_vrfy**), the action when the ACL is not defined is “deny”. This means that **acl\_smtp\_rcpt** must be defined in order to

receive any messages over an SMTP connection. For an example, see the ACL in the default configuration file.

### 43.15 Data for message ACLs

When a MAIL or RCPT ACL, or either of the DATA ACLs, is running, the variables that contain information about the host and the message's sender (for example, *\$sender\_host\_address* and *\$sender\_address*) are set, and can be used in ACL statements. In the case of RCPT (but not MAIL or DATA), *\$domain* and *\$local\_part* are set from the argument address. The entire SMTP command is available in *\$smtp\_command*.

When an ACL for the AUTH parameter of MAIL is running, the variables that contain information about the host are set, but *\$sender\_address* is not yet set. Section 33.2 contains a discussion of this parameter and how it is used.

The *\$message\_size* variable is set to the value of the SIZE parameter on the MAIL command at MAIL, RCPT and pre-data time, or to -1 if that parameter is not given. The value is updated to the true message size by the time the final DATA ACL is run (after the message data has been received).

The *\$rcpt\_count* variable increases by one for each RCPT command received. The *\$recipients\_count* variable increases by one each time a RCPT command is accepted, so while an ACL for RCPT is being processed, it contains the number of previously accepted recipients. At DATA time (for both the DATA ACLs), *\$rcpt\_count* contains the total number of RCPT commands, and *\$recipients\_count* contains the total number of accepted recipients.

### 43.16 Data for non-message ACLs

When an ACL is being run for AUTH, EHLO, ETRN, EXPN, HELO, STARTTLS, or VRFY, the remainder of the SMTP command line is placed in *\$smtp\_command\_argument*, and the entire SMTP command is available in *\$smtp\_command*. These variables can be tested using a **condition** condition. For example, here is an ACL for use with AUTH, which insists that either the session is encrypted, or the CRAM-MD5 authentication method is used. In other words, it does not permit authentication methods that use cleartext passwords on unencrypted connections.

```
acl_check_auth:
    accept encrypted = *
    accept condition = ${if eq{${uc:$smtp_command_argument}}\
                        {CRAM-MD5}}
    deny message = TLS encryption or CRAM-MD5 required
```

(Another way of applying this restriction is to arrange for the authenticators that use cleartext passwords not to be advertised when the connection is not encrypted. You can use the generic **server\_advertise\_condition** authenticator option to do this.)

### 43.17 Format of an ACL

An individual ACL consists of a number of statements. Each statement starts with a verb, optionally followed by a number of conditions and “modifiers”. Modifiers can change the way the verb operates, define error and log messages, set variables, insert delays, and vary the processing of accepted messages.

If all the conditions are met, the verb is obeyed. The same condition may be used (with different arguments) more than once in the same statement. This provides a means of specifying an “and” conjunction between conditions. For example:

```
deny dnslists = list1.example
             dnslists = list2.example
```

If there are no conditions, the verb is always obeyed. Exim stops evaluating the conditions and modifiers when it reaches a condition that fails. What happens then depends on the verb (and in one case, on a special modifier). Not all the conditions make sense at every testing point. For example, you cannot test a sender address in the ACL that is run for a VRFY command.



## 43.18 ACL verbs

The ACL verbs are as follows:

- **accept**: If all the conditions are met, the ACL returns “accept”. If any of the conditions are not met, what happens depends on whether **endpass** appears among the conditions (for syntax see below). If the failing condition is before **endpass**, control is passed to the next ACL statement; if it is after **endpass**, the ACL returns “deny”. Consider this statement, used to check a RCPT command:

```
accept domains = +local_domains
endpass
verify = recipient
```

If the recipient domain does not match the **domains** condition, control passes to the next statement. If it does match, the recipient is verified, and the command is accepted if verification succeeds. However, if verification fails, the ACL yields “deny”, because the failing condition is after **endpass**.

The **endpass** feature has turned out to be confusing to many people, so its use is not recommended nowadays. It is always possible to rewrite an ACL so that **endpass** is not needed, and it is no longer used in the default configuration.

If a **message** modifier appears on an **accept** statement, its action depends on whether or not **endpass** is present. In the absence of **endpass** (when an **accept** verb either accepts or passes control to the next statement), **message** can be used to vary the message that is sent when an SMTP command is accepted. For example, in a RCPT ACL you could have:

```
accept <some conditions>
message = OK, I will allow you through today
```

You can specify an SMTP response code, optionally followed by an “extended response code” at the start of the message, but the first digit must be the same as would be sent by default, which is 2 for an **accept** verb.

If **endpass** is present in an **accept** statement, **message** specifies an error message that is used when access is denied. This behaviour is retained for backward compatibility, but current “best practice” is to avoid the use of **endpass**.

- **defer**: If all the conditions are true, the ACL returns “defer” which, in an SMTP session, causes a 4xx response to be given. For a non-SMTP ACL, **defer** is the same as **deny**, because there is no way of sending a temporary error. For a RCPT command, **defer** is much the same as using a *redirect* router and `:defer:` while verifying, but the **defer** verb can be used in any ACL, and even for a recipient it might be a simpler approach.
- **deny**: If all the conditions are met, the ACL returns “deny”. If any of the conditions are not met, control is passed to the next ACL statement. For example,

```
deny dnslists = blackholes.mail-abuse.org
```

rejects commands from hosts that are on a DNS black list.

- **discard**: This verb behaves like **accept**, except that it returns “discard” from the ACL instead of “accept”. It is permitted only on ACLs that are concerned with receiving messages. When all the conditions are true, the sending entity receives a “success” response. However, **discard** causes recipients to be discarded. If it is used in an ACL for RCPT, just the one recipient is discarded; if used for MAIL, DATA or in the non-SMTP ACL, all the message’s recipients are discarded. Recipients that are discarded before DATA do not appear in the log line when the **received\_recipients** log selector is set.

If the **log\_message** modifier is set when **discard** operates, its contents are added to the line that is automatically written to the log. The **message** modifier operates exactly as it does for **accept**.

- **drop**: This verb behaves like **deny**, except that an SMTP connection is forcibly closed after the 5xx error message has been sent. For example:

```
drop    message    = I don't take more than 20 RCPTs
        condition = ${if > {${rcpt_count}}{20}}
```

There is no difference between **deny** and **drop** for the connect-time ACL. The connection is always dropped after sending a 550 response.

- **require**: If all the conditions are met, control is passed to the next ACL statement. If any of the conditions are not met, the ACL returns “deny”. For example, when checking a RCPT command,

```
require message = Sender did not verify
        verify  = sender
```

passes control to subsequent statements only if the message’s sender can be verified. Otherwise, it rejects the command. Note the positioning of the **message** modifier, before the **verify** condition. The reason for this is discussed in section 43.20.

- **warn**: If all the conditions are true, a line specified by the **log\_message** modifier is written to Exim’s main log. Control always passes to the next ACL statement. If any condition is false, the log line is not written. If an identical log line is requested several times in the same message, only one copy is actually written to the log. If you want to force duplicates to be written, use the **logwrite** modifier instead.

If **log\_message** is not present, a **warn** verb just checks its conditions and obeys any “immediate” modifiers (such as **control**, **set**, **logwrite**, **add\_header**, and **remove\_header**) that appear before the first failing condition. There is more about adding header lines in section 43.24.

If any condition on a **warn** statement cannot be completed (that is, there is some sort of defer), the log line specified by **log\_message** is not written. This does not include the case of a forced failure from a lookup, which is considered to be a successful completion. After a defer, no further conditions or modifiers in the **warn** statement are processed. The incident is logged, and the ACL continues to be processed, from the next statement onwards.

When one of the **warn** conditions is an address verification that fails, the text of the verification failure message is in *\$acl\_verify\_message*. If you want this logged, you must set it up explicitly. For example:

```
warn    !verify = sender
        log_message = sender verify failed: $acl_verify_message
```

At the end of each ACL there is an implicit unconditional **deny**.

As you can see from the examples above, the conditions and modifiers are written one to a line, with the first one on the same line as the verb, and subsequent ones on following lines. If you have a very long condition, you can continue it onto several physical lines by the usual backslash continuation mechanism. It is conventional to align the conditions vertically.

## 43.19 ACL variables

There are some special variables that can be set during ACL processing. They can be used to pass information between different ACLs, different invocations of the same ACL in the same SMTP connection, and between ACLs and the routers, transports, and filters that are used to deliver a message. The names of these variables must begin with *\$acl\_c* or *\$acl\_m*, followed either by a digit or an underscore, but the remainder of the name can be any sequence of alphanumeric characters and underscores that you choose. There is no limit on the number of ACL variables. The two sets act as follows:

- The values of those variables whose names begin with *\$acl\_c* persist throughout an SMTP connection. They are never reset. Thus, a value that is set while receiving one message is still available when receiving the next message on the same SMTP connection.
- The values of those variables whose names begin with *\$acl\_m* persist only while a message is being received. They are reset afterwards. They are also reset by MAIL, RSET, EHLO, HELO, and after starting up a TLS session.

When a message is accepted, the current values of all the ACL variables are preserved with the message and are subsequently made available at delivery time. The ACL variables are set by a modifier called **set**. For example:

```
accept hosts = whatever
    set acl_m4 = some value
accept authenticated = *
    set acl_c_auth = yes
```

**Note:** A leading dollar sign is not used when naming a variable that is to be set. If you want to set a variable without taking any action, you can use a **warn** verb without any other modifiers or conditions.

What happens if a syntactically valid but undefined ACL variable is referenced depends on the setting of the **strict\_acl\_vars** option. If it is false (the default), an empty string is substituted; if it is true, an error is generated.

Versions of Exim before 4.64 have a limited set of numbered variables, but their names are compatible, so there is no problem with upgrading.

## 43.20 Condition and modifier processing

An exclamation mark preceding a condition negates its result. For example:

```
deny    domains = *.dom.example
    !verify = recipient
```

causes the ACL to return “deny” if the recipient domain ends in *dom.example* and the recipient address cannot be verified. Sometimes negation can be used on the right-hand side of a condition. For example, these two statements are equivalent:

```
deny    hosts = !192.168.3.4
deny    !hosts = 192.168.3.4
```

However, for many conditions (**verify** being a good example), only left-hand side negation of the whole condition is possible.

The arguments of conditions and modifiers are expanded. A forced failure of an expansion causes a condition to be ignored, that is, it behaves as if the condition is true. Consider these two statements:

```
accept  senders = ${lookup{$host_name}lsearch\
    {/some/file}{$value}fail}
accept  senders = ${lookup{$host_name}lsearch\
    {/some/file}{$value}{} }
```

Each attempts to look up a list of acceptable senders. If the lookup succeeds, the returned list is searched, but if the lookup fails the behaviour is different in the two cases. The **fail** in the first statement causes the condition to be ignored, leaving no further conditions. The **accept** verb therefore succeeds. The second statement, however, generates an empty list when the lookup fails. No sender can match an empty list, so the condition fails, and therefore the **accept** also fails.

ACL modifiers appear mixed in with conditions in ACL statements. Some of them specify actions that are taken as the conditions for a statement are checked; others specify text for messages that are used when access is denied or a warning is generated. The **control** modifier affects the way an incoming message is handled.

The positioning of the modifiers in an ACL statement is important, because the processing of a verb ceases as soon as its outcome is known. Only those modifiers that have already been encountered will take effect. For example, consider this use of the **message** modifier:

```
require message = Can't verify sender
    verify = sender
    message = Can't verify recipient
    verify = recipient
    message = This message cannot be used
```

If sender verification fails, Exim knows that the result of the statement is “deny”, so it goes no further. The first **message** modifier has been seen, so its text is used as the error message. If sender verification succeeds, but recipient verification fails, the second message is used. If recipient verification succeeds, the third message becomes “current”, but is never used because there are no more conditions to cause failure.

For the **deny** verb, on the other hand, it is always the last **message** modifier that is used, because all the conditions must be true for rejection to happen. Specifying more than one **message** modifier does not make sense, and the message can even be specified after all the conditions. For example:

```
deny    hosts = ...
        !senders = *@my.domain.example
        message = Invalid sender from client host
```

The “deny” result does not happen until the end of the statement is reached, by which time Exim has set up the message.

## 43.21 ACL modifiers

The ACL modifiers are as follows:

### **add\_header** = <text>

This modifier specifies one or more header lines that are to be added to an incoming message, assuming, of course, that the message is ultimately accepted. For details, see section 43.24.

### **continue** = <text>

This modifier does nothing of itself, and processing of the ACL always continues with the next condition or modifier. The value of **continue** is in the side effects of expanding its argument. Typically this could be used to update a database. It is really just a syntactic tidiness, to avoid having to write rather ugly lines like this:

```
condition = ${if eq{0}{<some expansion>}{true}{true}}
```

Instead, all you need is

```
continue = <some expansion>
```

### **control** = <text>

This modifier affects the subsequent processing of the SMTP connection or of an incoming message that is accepted. The effect of the first type of control lasts for the duration of the connection, whereas the effect of the second type lasts only until the current message has been received. The message-specific controls always apply to the whole message, not to individual recipients, even if the **control** modifier appears in a RCPT ACL.

As there are now quite a few controls that can be applied, they are described separately in section 43.22. The **control** modifier can be used in several different ways. For example:

- It can be at the end of an **accept** statement:

```
accept  ...some conditions
        control = queue_only
```

In this case, the control is applied when this statement yields “accept”, in other words, when the conditions are all true.

- It can be in the middle of an **accept** statement:

```
accept  ...some conditions...
        control = queue_only
        ...some more conditions...
```

If the first set of conditions are true, the control is applied, even if the statement does not accept because one of the second set of conditions is false. In this case, some subsequent statement must yield “accept” for the control to be relevant.

- It can be used with **warn** to apply the control, leaving the decision about accepting or denying to a subsequent verb. For example:

```
warn      ...some conditions...
          control = freeze
accept    ...
```

This example of **warn** does not contain **message**, **log\_message**, or **logwrite**, so it does not add anything to the message and does not write a log entry.

- If you want to apply a control unconditionally, you can use it with a **require** verb. For example:

```
require   control = no_multiline_responses
```

#### **delay** = <time>

This modifier may appear in any ACL except notquit. It causes Exim to wait for the time interval before proceeding. However, when testing Exim using the **-bh** option, the delay is not actually imposed (an appropriate message is output instead). The time is given in the usual Exim notation, and the delay happens as soon as the modifier is processed. In an SMTP session, pending output is flushed before the delay is imposed.

Like **control**, **delay** can be used with **accept** or **deny**, for example:

```
deny      ...some conditions...
          delay = 30s
```

The delay happens if all the conditions are true, before the statement returns “deny”. Compare this with:

```
deny      delay = 30s
          ...some conditions...
```

which waits for 30s before processing the conditions. The **delay** modifier can also be used with **warn** and together with **control**:

```
warn      ...some conditions...
          delay = 2m
          control = freeze
accept    ...
```

If **delay** is encountered when the SMTP PIPELINING extension is in use, responses to several commands are no longer buffered and sent in one packet (as they would normally be) because all output is flushed before imposing the delay. This optimization is disabled so that a number of small delays do not appear to the client as one large aggregated delay that might provoke an unwanted timeout. You can, however, disable output flushing for **delay** by using a **control** modifier to set **no\_delay\_flush**.

#### **endpass**

This modifier, which has no argument, is recognized only in **accept** and **discard** statements. It marks the boundary between the conditions whose failure causes control to pass to the next statement, and the conditions whose failure causes the ACL to return “deny”. This concept has proved to be confusing to some people, so the use of **endpass** is no longer recommended as “best practice”. See the description of **accept** above for more details.

#### **log\_message** = <text>

This modifier sets up a message that is used as part of the log message if the ACL denies access or a **warn** statement’s conditions are true. For example:

```
require   log_message = wrong cipher suite $tls_in_cipher
          encrypted    = DES-CBC3-SHA
```

**log\_message** is also used when recipients are discarded by **discard**. For example:

```
discard   <some conditions>
          log_message = Discarded $local_part@$domain because...
```

When access is denied, **log\_message** adds to any underlying error message that may exist because of a condition failure. For example, while verifying a recipient address, a *:fail:* redirection might have already set up a message.

The message may be defined before the conditions to which it applies, because the string expansion does not happen until Exim decides that access is to be denied. This means that any variables that are set by the condition are available for inclusion in the message. For example, the *\$dnslist\_<xxx>* variables are set after a DNS black list lookup succeeds. If the expansion of **log\_message** fails, or if the result is an empty string, the modifier is ignored.

If you want to use a **warn** statement to log the result of an address verification, you can use *\$acl\_verify\_message* to include the verification error message.

If **log\_message** is used with a **warn** statement, “Warning:” is added to the start of the logged message. If the same warning log message is requested more than once while receiving a single email message, only one copy is actually logged. If you want to log multiple copies, use **logwrite** instead of **log\_message**. In the absence of **log\_message** and **logwrite**, nothing is logged for a successful **warn** statement.

If **log\_message** is not present and there is no underlying error message (for example, from the failure of address verification), but **message** is present, the **message** text is used for logging rejections. However, if any text for logging contains newlines, only the first line is logged. In the absence of both **log\_message** and **message**, a default built-in message is used for logging rejections.

**log\_reject\_target** = *<log name list>*

This modifier makes it possible to specify which logs are used for messages about ACL rejections. Its argument is a colon-separated list of words that can be “main”, “reject”, or “panic”. The default is *main:reject*. The list may be empty, in which case a rejection is not logged at all. For example, this ACL fragment writes no logging information when access is denied:

```
deny <some conditions>
    log_reject_target =
```

This modifier can be used in SMTP and non-SMTP ACLs. It applies to both permanent and temporary rejections. Its effect lasts for the rest of the current ACL.

**logwrite** = *<text>*

This modifier writes a message to a log file as soon as it is encountered when processing an ACL. (Compare **log\_message**, which, except in the case of **warn** and **discard**, is used only if the ACL statement denies access.) The **logwrite** modifier can be used to log special incidents in ACLs. For example:

```
accept <some special conditions>
    control    = freeze
    logwrite   = froze message because ...
```

By default, the message is written to the main log. However, it may begin with a colon, followed by a comma-separated list of log names, and then another colon, to specify exactly which logs are to be written. For example:

```
logwrite = :main,reject: text for main and reject logs
logwrite = :panic: text for panic log only
```

**message** = *<text>*

This modifier sets up a text string that is expanded and used as a response message when an ACL statement terminates the ACL with an “accept”, “deny”, or “defer” response. (In the case of the **accept** and **discard** verbs, there is some complication if **endpass** is involved; see the description of **accept** for details.)

The expansion of the message happens at the time Exim decides that the ACL is to end, not at the time it processes **message**. If the expansion fails, or generates an empty string, the modifier is ignored. Here is an example where **message** must be specified first, because the ACL ends with a rejection if the **hosts** condition fails:

```
require message = Host not recognized
hosts = 10.0.0.0/8
```

(Once a condition has failed, no further conditions or modifiers are processed.)

For ACLs that are triggered by SMTP commands, the message is returned as part of the SMTP response. The use of **message** with **accept** (or **discard**) is meaningful only for SMTP, as no message is returned when a non-SMTP message is accepted. In the case of the connect ACL, accepting with a message modifier overrides the value of **smtp\_banner**. For the EHLO/HELO ACL, a customized accept message may not contain more than one line (otherwise it will be truncated at the first newline and a panic logged), and it cannot affect the EHLO options.

When SMTP is involved, the message may begin with an overriding response code, consisting of three digits optionally followed by an “extended response code” of the form *n.n.n*, each code being followed by a space. For example:

```
deny message = 599 1.2.3 Host not welcome
hosts = 192.168.34.0/24
```

The first digit of the supplied response code must be the same as would be sent by default. A panic occurs if it is not. Exim uses a 550 code when it denies access, but for the predata ACL, note that the default success code is 354, not 2xx.

Notwithstanding the previous paragraph, for the QUIT ACL, unlike the others, the message modifier cannot override the 221 response code.

The text in a **message** modifier is literal; any quotes are taken as literals, but because the string is expanded, backslash escapes are processed anyway. If the message contains newlines, this gives rise to a multi-line SMTP response.

For ACLs that are called by an **acl** = ACL condition, the message is stored in *\$acl\_verify\_message*, from which the calling ACL may use it.

If **message** is used on a statement that verifies an address, the message specified overrides any message that is generated by the verification process. However, the original message is available in the variable *\$acl\_verify\_message*, so you can incorporate it into your message if you wish. In particular, if you want the text from **:fail:** items in *redirect* routers to be passed back as part of the SMTP response, you should either not use a **message** modifier, or make use of *\$acl\_verify\_message*.

For compatibility with previous releases of Exim, a **message** modifier that is used with a **warn** verb behaves in a similar way to the **add\_header** modifier, but this usage is now deprecated. However, **message** acts only when all the conditions are true, wherever it appears in an ACL command, whereas **add\_header** acts as soon as it is encountered. If **message** is used with **warn** in an ACL that is not concerned with receiving a message, it has no effect.

**queue** = <text>

This modifier specifies the use of a named queue for spool files for the message. It can only be used before the message is received (i.e. not in the DATA ACL). This could be used, for example, for known high-volume burst sources of traffic, or for quarantine of messages. Separate queue-runner processes will be needed for named queues. If the text after expansion is empty, the default queue is used.

**remove\_header** = <text>

This modifier specifies one or more header names in a colon-separated list that are to be removed from an incoming message, assuming, of course, that the message is ultimately accepted. For details, see section 43.25.

**set** <acl\_name> = <value>

This modifier puts a value into one of the ACL variables (see section 43.19).

**udpsend** = <parameters>

This modifier sends a UDP packet, for purposes such as statistics collection or behaviour monitoring. The parameters are expanded, and the result of the expansion must be a colon-separated list consisting of a destination server, port number, and the packet contents. The server can be speci-

fied as a host name or IPv4 or IPv6 address. The separator can be changed with the usual angle bracket syntax. For example, you might want to collect information on which hosts connect when:

```
udpsend = <; 2001:db8::dead:beef ; 1234 ;\
          $tod_zulu $sender_host_address
```

## 43.22 Use of the control modifier

The **control** modifier supports the following settings:

### **control = allow\_auth\_unadvertised**

This modifier allows a client host to use the SMTP AUTH command even when it has not been advertised in response to EHLO. Furthermore, because there are apparently some really broken clients that do this, Exim will accept AUTH after HELO (rather than EHLO) when this control is set. It should be used only if you really need it, and you should limit its use to those broken clients that do not work without it. For example:

```
warn hosts      = 192.168.34.25
  control = allow_auth_unadvertised
```

Normally, when an Exim server receives an AUTH command, it checks the name of the authentication mechanism that is given in the command to ensure that it matches an advertised mechanism. When this control is set, the check that a mechanism has been advertised is bypassed. Any configured mechanism can be used by the client. This control is permitted only in the connection and HELO ACLs.

### **control = caseful\_local\_part**

### **control = caselower\_local\_part**

These two controls are permitted only in the ACL specified by **acl\_smtp\_rcpt** (that is, during RCPT processing). By default, the contents of *\$local\_part* are lower cased before ACL processing. If “caseful\_local\_part” is specified, any uppercase letters in the original local part are restored in *\$local\_part* for the rest of the ACL, or until a control that sets “caselower\_local\_part” is encountered.

These controls affect only the current recipient. Moreover, they apply only to local part handling that takes place directly in the ACL (for example, as a key in lookups). If a test to verify the recipient is obeyed, the case-related handling of the local part during the verification is controlled by the router configuration (see the **caseful\_local\_part** generic router option).

This facility could be used, for example, to add a spam score to local parts containing upper case letters. For example, using *\$acl\_m4* to accumulate the spam score:

```
warn  control = caseful_local_part
      set acl_m4 = ${eval:\
                    $acl_m4 + \
                    ${if match{$local_part}{[A-Z]}}{1}{0}}\
                    }
      control = caselower_local_part
```

Notice that we put back the lower cased version afterwards, assuming that is what is wanted for subsequent tests.

### **control = cutthrough\_delivery/<options>**

This option requests delivery be attempted while the item is being received.

The option is usable in the RCPT ACL. If enabled for a message received via smtp and routed to an smtp transport, and only one transport, interface, destination host and port combination is used for all recipients of the message, then the delivery connection is made while the receiving connection is open and data is copied from one to the other.

An attempt to set this option for any recipient but the first for a mail will be quietly ignored. If a recipient-verify callout (with *use\_sender*) connection is subsequently requested in the same ACL it is held open and used for any subsequent recipients and the data, otherwise one is made after the initial RCPT ACL completes.



Note that routers are used in verify mode, and cannot depend on content of received headers. Note also that headers cannot be modified by any of the post-data ACLs (DATA, MIME and DKIM). Headers may be modified by routers (subject to the above) and transports.

The Received-By: header is generated as soon as the body reception starts, rather than the traditional time after the full message is received; this will affect the timestamp.

All the usual ACLs are called; if one results in the message being rejected, all effort spent in delivery (including the costs on the ultimate destination) will be wasted. Note that in the case of data-time ACLs this includes the entire message body.

Cutthrough delivery is not supported via transport-filters or when DKIM signing of outgoing messages is done, because it sends data to the ultimate destination before the entire message has been received from the source. It is not supported for messages received with the SMTP PRDR or CHUNKING options in use.

Should the ultimate destination system positively accept or reject the mail, a corresponding indication is given to the source system and nothing is queued. If the item is successfully delivered in cutthrough mode the delivery log lines are tagged with ">>" rather than "=>" and appear before the acceptance "<=" line.

If there is a temporary error the item is queued for later delivery in the usual fashion. This behaviour can be adjusted by appending the option **defer=<value>** to the control; the default value is "spool" and the alternate value "pass" copies an SMTP defer response from the target back to the initiator and does not queue the message. Note that this is independent of any recipient verify conditions in the ACL.

Delivery in this mode avoids the generation of a bounce mail to a (possibly faked) sender when the destination system is doing content-scan based rejection.

#### **control = debug/<options>**

This control turns on debug logging, almost as though Exim had been invoked with `-d`, with the output going to a new logfile in the usual logs directory, by default called *debuglog*. The filename can be adjusted with the *tag* option, which may access any variables already defined. The logging may be adjusted with the *opts* option, which takes the same values as the `-d` command-line option. Logging started this way may be stopped, and the file removed, with the *kill* option. Some examples (which depend on variables that don't exist in all contexts):

```
control = debug
control = debug/tag=.$sender_host_address
control = debug/opts=+expand+acl
control = debug/tag=.$message_exim_id/opts=+expand
control = debug/kill
```

#### **control = dkim\_disable\_verify**

This control turns off DKIM verification processing entirely. For details on the operation and configuration of DKIM, see chapter 57.

#### **control = dscp/<value>**

This option causes the DSCP value associated with the socket for the inbound connection to be adjusted to a given value, given as one of a number of fixed strings or to numeric value. The **-bl:dscp** option may be used to ask Exim which names it knows of. Common values include *throughput*, *mincost*, and on newer systems *ef*, *af41*, etc. Numeric values may be in the range 0 to 0x3F.

The outbound packets from Exim will be marked with this value in the header (for IPv4, the TOS field; for IPv6, the TCLASS field); there is no guarantee that these values will have any effect, not be stripped by networking equipment, or do much of anything without cooperation with your Network Engineer and those of all network operators between the source and destination.

**control = enforce\_sync**

**control = no\_enforce\_sync**

These controls make it possible to be selective about when SMTP synchronization is enforced. The global option **smtp\_enforce\_sync** specifies the initial state of the switch (it is true by default). See the description of this option in chapter 14 for details of SMTP synchronization checking.

The effect of these two controls lasts for the remainder of the SMTP connection. They can appear in any ACL except the one for the non-SMTP messages. The most straightforward place to put them is in the ACL defined by **acl\_smtp\_connect**, which is run at the start of an incoming SMTP connection, before the first synchronization check. The expected use is to turn off the synchronization checks for badly-behaved hosts that you nevertheless need to work with.

**control = fakedefer/<message>**

This control works in exactly the same way as **fakereject** (described below) except that it causes an SMTP 450 response after the message data instead of a 550 response. You must take care when using **fakedefer** because it causes the messages to be duplicated when the sender retries. Therefore, you should not use **fakedefer** if the message is to be delivered normally.

**control = fakereject/<message>**

This control is permitted only for the MAIL, RCPT, and DATA ACLs, in other words, only when an SMTP message is being received. If Exim accepts the message, instead the final 250 response, a 550 rejection message is sent. However, Exim proceeds to deliver the message as normal. The control applies only to the current message, not to any subsequent ones that may be received in the same SMTP connection.

The text for the 550 response is taken from the **control** modifier. If no message is supplied, the following is used:

```
550-Your message has been rejected but is being
550-kept for evaluation.
550-If it was a legitimate message, it may still be
550 delivered to the target recipient(s).
```

This facility should be used with extreme caution.

**control = freeze**

This control is permitted only for the MAIL, RCPT, DATA, and non-SMTP ACLs, in other words, only when a message is being received. If the message is accepted, it is placed on Exim's queue and frozen. The control applies only to the current message, not to any subsequent ones that may be received in the same SMTP connection.

This modifier can optionally be followed by **/no\_tell**. If the global option **freeze\_tell** is set, it is ignored for the current message (that is, nobody is told about the freezing), provided all the **control=freeze** modifiers that are obeyed for the current message have the **/no\_tell** option.

**control = no\_delay\_flush**

Exim normally flushes SMTP output before implementing a delay in an ACL, to avoid unexpected timeouts in clients when the SMTP PIPELINING extension is in use. This control, as long as it is encountered before the **delay** modifier, disables such output flushing.

**control = no\_callout\_flush**

Exim normally flushes SMTP output before performing a callout in an ACL, to avoid unexpected timeouts in clients when the SMTP PIPELINING extension is in use. This control, as long as it is encountered before the **verify** condition that causes the callout, disables such output flushing.

**control = no\_mbox\_unspool**

This control is available when Exim is compiled with the content scanning extension. Content scanning may require a copy of the current message, or parts of it, to be written in "mbox format" to a spool file, for passing to a virus or spam scanner. Normally, such copies are deleted when they are no longer needed. If this control is set, the copies are not deleted. The control applies only to the current message, not to any subsequent ones that may be received in the same SMTP connection. It is provided for debugging purposes and is unlikely to be useful in production.

### **control = no\_multiline\_responses**

This control is permitted for any ACL except the one for non-SMTP messages. It seems that there are broken clients in use that cannot handle multiline SMTP responses, despite the fact that RFC 821 defined them over 20 years ago.

If this control is set, multiline SMTP responses from ACL rejections are suppressed. One way of doing this would have been to put out these responses as one long line. However, RFC 2821 specifies a maximum of 512 bytes per response (“use multiline responses for more” it says – ha!), and some of the responses might get close to that. So this facility, which is after all only a sop to broken clients, is implemented by doing two very easy things:

- Extra information that is normally output as part of a rejection caused by sender verification failure is omitted. Only the final line (typically “sender verification failed”) is sent.
- If a **message** modifier supplies a multiline response, only the first line is output.

The setting of the switch can, of course, be made conditional on the calling host. Its effect lasts until the end of the SMTP connection.

### **control = no\_pipelining**

This control turns off the advertising of the PIPELINING extension to SMTP in the current session. To be useful, it must be obeyed before Exim sends its response to an EHLO command. Therefore, it should normally appear in an ACL controlled by **acl\_smtp\_connect** or **acl\_smtp\_helo**. See also **pipelining\_advertise\_hosts**.

### **control = queue\_only**

This control is permitted only for the MAIL, RCPT, DATA, and non-SMTP ACLs, in other words, only when a message is being received. If the message is accepted, it is placed on Exim’s queue and left there for delivery by a subsequent queue runner. No immediate delivery process is started. In other words, it has the effect as the **queue\_only** global option. However, the control applies only to the current message, not to any subsequent ones that may be received in the same SMTP connection.

### **control = submission/<options>**

This control is permitted only for the MAIL, RCPT, and start of data ACLs (the latter is the one defined by **acl\_smtp\_predata**). Setting it tells Exim that the current message is a submission from a local MUA. In this case, Exim operates in “submission mode”, and applies certain fixups to the message if necessary. For example, it adds a *Date:* header line if one is not present. This control is not permitted in the **acl\_smtp\_data** ACL, because that is too late (the message has already been created).

Chapter 47 describes the processing that Exim applies to messages. Section 47.1 covers the processing that happens in submission mode; the available options for this control are described there. The control applies only to the current message, not to any subsequent ones that may be received in the same SMTP connection.

### **control = suppress\_local\_fixups**

This control applies to locally submitted (non TCP/IP) messages, and is the complement of **control = submission**. It disables the fixups that are normally applied to locally-submitted messages. Specifically:

- Any *Sender:* header line is left alone (in this respect, it is a dynamic version of **local\_sender\_retain**).
- No *Message-ID:*, *From:*, or *Date:* header lines are added.
- There is no check that *From:* corresponds to the actual sender.

This control may be useful when a remotely-originated message is accepted, passed to some scanning program, and then re-submitted for delivery. It can be used only in the **acl\_smtp\_mail**, **acl\_smtp\_rcpt**, **acl\_smtp\_predata**, and **acl\_not\_smtp\_start** ACLs, because it has to be set before the message’s data is read.

**Note:** This control applies only to the current message, not to any others that are being submitted at the same time using **-bs** or **-bS**.

### **control = utf8\_downconvert**

This control enables conversion of UTF-8 in message addresses to a-label form. For details see section 59.1.

## **43.23 Summary of message fixup control**

All four possibilities for message fixups can be specified:

- Locally submitted, fixups applied: the default.
- Locally submitted, no fixups applied: use `control = suppress_local_fixups`.
- Remotely submitted, no fixups applied: the default.
- Remotely submitted, fixups applied: use `control = submission`.

## **43.24 Adding header lines in ACLs**

The **add\_header** modifier can be used to add one or more extra header lines to an incoming message, as in this example:

```
warn dnslists = sbl.spamhaus.org : \  
                dialup.mail-abuse.org  
add_header = X-blacklisted-at: $dnslist_domain
```

The **add\_header** modifier is permitted in the MAIL, RCPT, PREDATA, DATA, MIME, DKIM, and non-SMTP ACLs (in other words, those that are concerned with receiving a message). The message must ultimately be accepted for **add\_header** to have any significant effect. You can use **add\_header** with any ACL verb, including **deny** (though this is potentially useful only in a RCPT ACL).

Headers will not be added to the message if the modifier is used in DATA, MIME or DKIM ACLs for a message delivered by cutthrough routing.

Leading and trailing newlines are removed from the data for the **add\_header** modifier; if it then contains one or more newlines that are not followed by a space or a tab, it is assumed to contain multiple header lines. Each one is checked for valid syntax; X-ACL-Warn: is added to the front of any line that is not a valid header line.

Added header lines are accumulated during the MAIL, RCPT, and predata ACLs. They are added to the message before processing the DATA and MIME ACLs. However, if an identical header line is requested more than once, only one copy is actually added to the message. Further header lines may be accumulated during the DATA and MIME ACLs, after which they are added to the message, again with duplicates suppressed. Thus, it is possible to add two identical header lines to an SMTP message, but only if one is added before DATA and one after. In the case of non-SMTP messages, new headers are accumulated during the non-SMTP ACLs, and are added to the message after all the ACLs have run. If a message is rejected after DATA or by the non-SMTP ACL, all added header lines are included in the entry that is written to the reject log.

Header lines are not visible in string expansions of message headers until they are added to the message. It follows that header lines defined in the MAIL, RCPT, and predata ACLs are not visible until the DATA ACL and MIME ACLs are run. Similarly, header lines that are added by the DATA or MIME ACLs are not visible in those ACLs. Because of this restriction, you cannot use header lines as a way of passing data between (for example) the MAIL and RCPT ACLs. If you want to do this, you can use ACL variables, as described in section 43.19.

The list of headers yet to be added is given by the **\$headers\_added** variable.

The **add\_header** modifier acts immediately as it is encountered during the processing of an ACL. Notice the difference between these two cases:

```
accept add_header = ADDED: some text  
    <some condition>
```

```
accept <some condition>
    add_header = ADDED: some text
```

In the first case, the header line is always added, whether or not the condition is true. In the second case, the header line is added only if the condition is true. Multiple occurrences of **add\_header** may occur in the same ACL statement. All those that are encountered before a condition fails are honoured.

For compatibility with previous versions of Exim, a **message** modifier for a **warn** verb acts in the same way as **add\_header**, except that it takes effect only if all the conditions are true, even if it appears before some of them. Furthermore, only the last occurrence of **message** is honoured. This usage of **message** is now deprecated. If both **add\_header** and **message** are present on a **warn** verb, both are processed according to their specifications.

By default, new header lines are added to a message at the end of the existing header lines. However, you can specify that any particular header line should be added right at the start (before all the *Received:* lines), immediately after the first block of *Received:* lines, or immediately before any line that is not a *Received:* or *Resent-something:* header.

This is done by specifying “:at\_start:”, “:after\_received:”, or “:at\_start\_rfc:” (or, for completeness, “:at\_end:”) before the text of the header line, respectively. (Header text cannot start with a colon, as there has to be a header name first.) For example:

```
warn add_header = \
    :after_received:X-My-Header: something or other...
```

If more than one header line is supplied in a single **add\_header** modifier, each one is treated independently and can therefore be placed differently. If you add more than one line at the start, or after the *Received:* block, they end up in reverse order.

**Warning:** This facility currently applies only to header lines that are added in an ACL. It does NOT work for header lines that are added in a system filter or in a router or transport.

## 43.25 Removing header lines in ACLs

The **remove\_header** modifier can be used to remove one or more header lines from an incoming message, as in this example:

```
warn    message          = Remove internal headers
    remove_header = x-route-mail1 : x-route-mail2
```

The **remove\_header** modifier is permitted in the MAIL, RCPT, PREDATA, DATA, MIME, DKIM, and non-SMTP ACLs (in other words, those that are concerned with receiving a message). The message must ultimately be accepted for **remove\_header** to have any significant effect. You can use **remove\_header** with any ACL verb, including **deny**, though this is really not useful for any verb that doesn’t result in a delivered message.

Headers will not be removed from the message if the modifier is used in DATA, MIME or DKIM ACLs for a message delivered by cutthrough routing.

More than one header can be removed at the same time by using a colon separated list of header names. The header matching is case insensitive. Wildcards are not permitted, nor is list expansion performed, so you cannot use hostlists to create a list of headers, however both connection and message variable expansion are performed (**\$acl\_c\_\*** and **\$acl\_m\_\***), illustrated in this example:

```
warn    hosts            = +internal_hosts
    set acl_c_ihdrs = x-route-mail1 : x-route-mail2
warn    message          = Remove internal headers
    remove_header = $acl_c_ihdrs
```

Removed header lines are accumulated during the MAIL, RCPT, and predata ACLs. They are removed from the message before processing the DATA and MIME ACLs. There is no harm in attempting to remove the same header twice nor is removing a non-existent header. Further header lines to be removed may be accumulated during the DATA and MIME ACLs, after which they are

removed from the message, if present. In the case of non-SMTP messages, headers to be removed are accumulated during the non-SMTP ACLs, and are removed from the message after all the ACLs have run. If a message is rejected after DATA or by the non-SMTP ACL, there really is no effect because there is no logging of what headers would have been removed.

Header lines are not visible in string expansions until the DATA phase when it is received. Any header lines removed in the MAIL, RCPT, and predata ACLs are not visible in the DATA ACL and MIME ACLs. Similarly, header lines that are removed by the DATA or MIME ACLs are still visible in those ACLs. Because of this restriction, you cannot use header lines as a way of controlling data passed between (for example) the MAIL and RCPT ACLs. If you want to do this, you should instead use ACL variables, as described in section 43.19.

The **remove\_header** modifier acts immediately as it is encountered during the processing of an ACL. Notice the difference between these two cases:

```
accept remove_header = X-Internal
    <some condition>

accept <some condition>
    remove_header = X-Internal
```

In the first case, the header line is always removed, whether or not the condition is true. In the second case, the header line is removed only if the condition is true. Multiple occurrences of **remove\_header** may occur in the same ACL statement. All those that are encountered before a condition fails are honoured.

**Warning:** This facility currently applies only to header lines that are present during ACL processing. It does NOT remove header lines that are added in a system filter or in a router or transport.

## 43.26 ACL conditions

Some of the conditions listed in this section are available only when Exim is compiled with the content-scanning extension. They are included here briefly for completeness. More detailed descriptions can be found in the discussion on content scanning in chapter 44.

Not all conditions are relevant in all circumstances. For example, testing senders and recipients does not make sense in an ACL that is being run as the result of the arrival of an ETRN command, and checks on message headers can be done only in the ACLs specified by **acl\_smtp\_data** and **acl\_not\_smtp**. You can use the same condition (with different parameters) more than once in the same ACL statement. This provides a way of specifying an “and” conjunction. The conditions are as follows:

**acl** = <name of acl or ACL string or file name >

The possible values of the argument are the same as for the **acl\_smtp\_xxx** options. The named or inline ACL is run. If it returns “accept” the condition is true; if it returns “deny” the condition is false. If it returns “defer”, the current ACL returns “defer” unless the condition is on a **warn** verb. In that case, a “defer” return makes the condition false. This means that further processing of the **warn** verb ceases, but processing of the ACL continues.

If the argument is a named ACL, up to nine space-separated optional values can be appended; they appear within the called ACL in \$acl\_arg1 to \$acl\_arg9, and \$acl\_narg is set to the count of values. Previous values of these variables are restored after the call returns. The name and values are expanded separately. Note that spaces in complex expansions which are used as arguments will act as argument separators.

If the nested **acl** returns “drop” and the outer condition denies access, the connection is dropped. If it returns “discard”, the verb must be **accept** or **discard**, and the action is taken immediately – no further conditions are tested.

ACLs may be nested up to 20 deep; the limit exists purely to catch runaway loops. This condition allows you to use different ACLs in different circumstances. For example, different ACLs can be used to handle RCPT commands for different local users or different local domains.

**authenticated** = <*string list*>

If the SMTP connection is not authenticated, the condition is false. Otherwise, the name of the authenticator is tested against the list. To test for authentication by any authenticator, you can set

```
authenticated = *
```

**condition** = <*string*>

This feature allows you to make up custom conditions. If the result of expanding the string is an empty string, the number zero, or one of the strings “no” or “false”, the condition is false. If the result is any non-zero number, or one of the strings “yes” or “true”, the condition is true. For any other value, some error is assumed to have occurred, and the ACL returns “defer”. However, if the expansion is forced to fail, the condition is ignored. The effect is to treat it as true, whether it is positive or negative.

**decode** = <*location*>

This condition is available only when Exim is compiled with the content-scanning extension, and it is allowed only in the ACL defined by **acl\_smtp\_mime**. It causes the current MIME part to be decoded into a file. If all goes well, the condition is true. It is false only if there are problems such as a syntax error or a memory shortage. For more details, see chapter 44.

**dnslists** = <*list of domain names and other data*>

This condition checks for entries in DNS black lists. These are also known as “RBL lists”, after the original Realtime Blackhole List, but note that the use of the lists at *mail-abuse.org* now carries a charge. There are too many different variants of this condition to describe briefly here. See sections 43.27–43.37 for details.

**domains** = <*domain list*>

This condition is relevant only after a RCPT command. It checks that the domain of the recipient address is in the domain list. If percent-hack processing is enabled, it is done before this test is done. If the check succeeds with a lookup, the result of the lookup is placed in *\$domain\_data* until the next **domains** test.

**Note carefully** (because many people seem to fall foul of this): you cannot use **domains** in a DATA ACL.

**encrypted** = <*string list*>

If the SMTP connection is not encrypted, the condition is false. Otherwise, the name of the cipher suite in use is tested against the list. To test for encryption without testing for any specific cipher suite(s), set

```
encrypted = *
```

**hosts** = <*host list*>

This condition tests that the calling host matches the host list. If you have name lookups or wildcarded host names and IP addresses in the same host list, you should normally put the IP addresses first. For example, you could have:

```
accept hosts = 10.9.8.7 : dbm;/etc/friendly/hosts
```

The lookup in this example uses the host name for its key. This is implied by the lookup type “dbm”. (For a host address lookup you would use “net-dbm” and it wouldn’t matter which way round you had these two items.)

The reason for the problem with host names lies in the left-to-right way that Exim processes lists. It can test IP addresses without doing any DNS lookups, but when it reaches an item that requires a host name, it fails if it cannot find a host name to compare with the pattern. If the above list is given in the opposite order, the **accept** statement fails for a host whose name cannot be found, even if its IP address is 10.9.8.7.

If you really do want to do the name check first, and still recognize the IP address even if the name lookup fails, you can rewrite the ACL like this:

```
accept hosts = dbm;/etc/friendly/hosts
accept hosts = 10.9.8.7
```

The default action on failing to find the host name is to assume that the host is not in the list, so the first **accept** statement fails. The second statement can then check the IP address.

If a **hosts** condition is satisfied by means of a lookup, the result of the lookup is made available in the *\$host\_data* variable. This allows you, for example, to set up a statement like this:

```
deny hosts = net-lsearch;/some/file
message = $host_data
```

which gives a custom error message for each denied host.

**local\_parts** = *<local part list>*

This condition is relevant only after a RCPT command. It checks that the local part of the recipient address is in the list. If percent-hack processing is enabled, it is done before this test. If the check succeeds with a lookup, the result of the lookup is placed in *\$local\_part\_data*, which remains set until the next **local\_parts** test.

**malware** = *<option>*

This condition is available only when Exim is compiled with the content-scanning extension. It causes the incoming message to be scanned for viruses. For details, see chapter 44.

**mime\_regex** = *<list of regular expressions>*

This condition is available only when Exim is compiled with the content-scanning extension, and it is allowed only in the ACL defined by **acl\_smtp\_mime**. It causes the current MIME part to be scanned for a match with any of the regular expressions. For details, see chapter 44.

**ratelimit** = *<parameters>*

This condition can be used to limit the rate at which a user or host submits messages. Details are given in section 43.38.

**recipients** = *<address list>*

This condition is relevant only after a RCPT command. It checks the entire recipient address against a list of recipients.

**regex** = *<list of regular expressions>*

This condition is available only when Exim is compiled with the content-scanning extension, and is available only in the DATA, MIME, and non-SMTP ACLs. It causes the incoming message to be scanned for a match with any of the regular expressions. For details, see chapter 44.

**sender\_domains** = *<domain list>*

This condition tests the domain of the sender of the message against the given domain list. **Note:** The domain of the sender address is in *\$sender\_address\_domain*. It is *not* put in *\$domain* during the testing of this condition. This is an exception to the general rule for testing domain lists. It is done this way so that, if this condition is used in an ACL for a RCPT command, the recipient's domain (which is in *\$domain*) can be used to influence the sender checking.

**Warning:** It is a bad idea to use this condition on its own as a control on relaying, because sender addresses are easily, and commonly, forged.

**senders** = *<address list>*

This condition tests the sender of the message against the given list. To test for a bounce message, which has an empty sender, set

```
senders = :
```

**Warning:** It is a bad idea to use this condition on its own as a control on relaying, because sender addresses are easily, and commonly, forged.

**spam** = *<username>*

This condition is available only when Exim is compiled with the content-scanning extension. It causes the incoming message to be scanned by SpamAssassin. For details, see chapter 44.

**verify** = **certificate**

This condition is true in an SMTP session if the session is encrypted, and a certificate was received from the client, and the certificate was verified. The server requests a certificate only if the client matches **tls\_verify\_hosts** or **tls\_try\_verify\_hosts** (see chapter 42).



### **verify = csa**

This condition checks whether the sending host (the client) is authorized to send email. Details of how this works are given in section 43.50.

### **verify = header\_names\_ascii**

This condition is relevant only in an ACL that is run after a message has been received, that is, in an ACL specified by **acl\_smtp\_data** or **acl\_not\_smtp**. It checks all header names (not the content) to make sure there are no non-ASCII characters, also excluding control characters. The allowable characters are decimal ASCII values 33 through 126.

Exim itself will handle headers with non-ASCII characters, but it can cause problems for downstream applications, so this option will allow their detection and rejection in the DATA ACL's.

### **verify = header\_sender/<options>**

This condition is relevant only in an ACL that is run after a message has been received, that is, in an ACL specified by **acl\_smtp\_data** or **acl\_not\_smtp**. It checks that there is a verifiable address in at least one of the *Sender:*, *Reply-To:*, or *From:* header lines. Such an address is loosely thought of as a “sender” address (hence the name of the test). However, an address that appears in one of these headers need not be an address that accepts bounce messages; only sender addresses in envelopes are required to accept bounces. Therefore, if you use the callout option on this check, you might want to arrange for a non-empty address in the MAIL command.

Details of address verification and the options are given later, starting at section 43.44 (callouts are described in section 43.45). You can combine this condition with the **senders** condition to restrict it to bounce messages only:

```
deny      senders = :
          message = A valid sender header is required for bounces
          !verify = header_sender
```

### **verify = header\_syntax**

This condition is relevant only in an ACL that is run after a message has been received, that is, in an ACL specified by **acl\_smtp\_data** or **acl\_not\_smtp**. It checks the syntax of all header lines that can contain lists of addresses (*Sender:*, *From:*, *Reply-To:*, *To:*, *Cc:*, and *Bcc:*), returning true if there are no problems. Unqualified addresses (local parts without domains) are permitted only in locally generated messages and from hosts that match **sender\_unqualified\_hosts** or **recipient\_unqualified\_hosts**, as appropriate.

Note that this condition is a syntax check only. However, a common spamming ploy used to be to send syntactically invalid headers such as

```
To: @
```

and this condition can be used to reject such messages, though they are not as common as they used to be.

### **verify = helo**

This condition is true if a HELO or EHLO command has been received from the client host, and its contents have been verified. If there has been no previous attempt to verify the HELO/EHLO contents, it is carried out when this condition is encountered. See the description of the **helo\_verify\_hosts** and **helo\_try\_verify\_hosts** options for details of how to request verification independently of this condition.

For SMTP input that does not come over TCP/IP (the **-bs** command line option), this condition is always true.

### **verify = not\_blind**

This condition checks that there are no blind (bcc) recipients in the message. Every envelope recipient must appear either in a *To:* header line or in a *Cc:* header line for this condition to be true. Local parts are checked case-sensitively; domains are checked case-insensitively. If *Resent-To:* or *Resent-Cc:* header lines exist, they are also checked. This condition can be used only in a DATA or non-SMTP ACL.

There are, of course, many legitimate messages that make use of blind (bcc) recipients. This check should not be used on its own for blocking messages.

**verify = recipient/<options>**

This condition is relevant only after a RCPT command. It verifies the current recipient. Details of address verification are given later, starting at section 43.44. After a recipient has been verified, the value of *\$address\_data* is the last value that was set while routing the address. This applies even if the verification fails. When an address that is being verified is redirected to a single address, verification continues with the new address, and in that case, the subsequent value of *\$address\_data* is the value for the child address.

**verify = reverse\_host\_lookup/<options>**

This condition ensures that a verified host name has been looked up from the IP address of the client host. (This may have happened already if the host name was needed for checking a host list, or if the host matched **host\_lookup**.) Verification ensures that the host name obtained from a reverse DNS lookup, or one of its aliases, does, when it is itself looked up in the DNS, yield the original IP address.

There is one possible option, *defer\_ok*. If this is present and a DNS operation returns a temporary error, the verify condition succeeds.

If this condition is used for a locally generated message (that is, when there is no client host involved), it always succeeds.

**verify = sender/<options>**

This condition is relevant only after a MAIL or RCPT command, or after a message has been received (the **acl\_smtp\_data** or **acl\_not\_smtp** ACLs). If the message's sender is empty (that is, this is a bounce message), the condition is true. Otherwise, the sender address is verified.

If there is data in the *\$address\_data* variable at the end of routing, its value is placed in *\$sender\_address\_data* at the end of verification. This value can be used in subsequent conditions and modifiers in the same ACL statement. It does not persist after the end of the current statement. If you want to preserve the value for longer, you can save it in an ACL variable.

Details of verification are given later, starting at section 43.44. Exim caches the result of sender verification, to avoid doing it more than once per message.

**verify = sender=<address>/<options>**

This is a variation of the previous option, in which a modified address is verified as a sender.

Note that '/' is legal in local-parts; if the address may have such (eg. is generated from the received message) they must be protected from the options parsing by doubling:

```
verify = sender=${sg{${address:$h_sender:}}{/}{//}}
```

## 43.27 Using DNS lists

In its simplest form, the **dnslists** condition tests whether the calling host is on at least one of a number of DNS lists by looking up the inverted IP address in one or more DNS domains. (Note that DNS list domains are not mail domains, so the + syntax for named lists doesn't work - it is used for special options instead.) For example, if the calling host's IP address is 192.168.62.43, and the ACL statement is

```
deny dnslists = blackholes.mail-abuse.org : \
                dialups.mail-abuse.org
```

the following records are looked up:

```
43.62.168.192.blackholes.mail-abuse.org
43.62.168.192.dialups.mail-abuse.org
```

As soon as Exim finds an existing DNS record, processing of the list stops. Thus, multiple entries on the list provide an "or" conjunction. If you want to test that a host is on more than one list (an "and" conjunction), you can use two separate conditions:

```
deny dnslists = blackholes.mail-abuse.org
dnslists = dialups.mail-abuse.org
```

If a DNS lookup times out or otherwise fails to give a decisive answer, Exim behaves as if the host does not match the list item, that is, as if the DNS record does not exist. If there are further items in the DNS list, they are processed.

This is usually the required action when **dnslists** is used with **deny** (which is the most common usage), because it prevents a DNS failure from blocking mail. However, you can change this behaviour by putting one of the following special items in the list:

```
+include_unknown    behave as if the item is on the list
+exclude_unknown    behave as if the item is not on the list (default)
+defer_unknown      give a temporary error
```

Each of these applies to any subsequent items on the list. For example:

```
deny dnslists = +defer_unknown : foo.bar.example
```

Testing the list of domains stops as soon as a match is found. If you want to warn for one list and block for another, you can use two different statements:

```
deny dnslists = blackholes.mail-abuse.org
warn message = X-Warn: sending host is on dialups list
dnslists = dialups.mail-abuse.org
```

DNS list lookups are cached by Exim for the duration of the SMTP session (but limited by the DNS return TTL value), so a lookup based on the IP address is done at most once for any incoming connection (assuming long-enough TTL). Exim does not share information between multiple incoming connections (but your local name server cache should be active).

## 43.28 Specifying the IP address for a DNS list lookup

By default, the IP address that is used in a DNS list lookup is the IP address of the calling host. However, you can specify another IP address by listing it after the domain name, introduced by a slash. For example:

```
deny dnslists = black.list.tld/192.168.1.2
```

This feature is not very helpful with explicit IP addresses; it is intended for use with IP addresses that are looked up, for example, the IP addresses of the MX hosts or nameservers of an email sender address. For an example, see section 43.30 below.

## 43.29 DNS lists keyed on domain names

There are some lists that are keyed on domain names rather than inverted IP addresses (see for example the *domain based zones* link at <http://www.rfc-ignorant.org/>). No reversing of components is used with these lists. You can change the name that is looked up in a DNS list by listing it after the domain name, introduced by a slash. For example,

```
deny message = Sender's domain is listed at $dnslist_domain
dnslists = dsn.rfc-ignorant.org/$sender_address_domain
```

This particular example is useful only in ACLs that are obeyed after the RCPT or DATA commands, when a sender address is available. If (for example) the message's sender is *user@tld.example* the name that is looked up by this example is

```
tld.example.dsn.rfc-ignorant.org
```

A single **dnslists** condition can contain entries for both names and IP addresses. For example:

```
deny dnslists = sbl.spamhaus.org : \
dsn.rfc-ignorant.org/$sender_address_domain
```

The first item checks the sending host's IP address; the second checks a domain name. The whole condition is true if either of the DNS lookups succeeds.

### 43.30 Multiple explicit keys for a DNS list

The syntax described above for looking up explicitly-defined values (either names or IP addresses) in a DNS blacklist is a simplification. After the domain name for the DNS list, what follows the slash can in fact be a list of items. As with all lists in Exim, the default separator is a colon. However, because this is a sublist within the list of DNS blacklist domains, it is necessary either to double the separators like this:

```
dnslists = black.list.tld/name.1::name.2
```

or to change the separator character, like this:

```
dnslists = black.list.tld/<;name.1;name.2
```

If an item in the list is an IP address, it is inverted before the DNS blacklist domain is appended. If it is not an IP address, no inversion occurs. Consider this condition:

```
dnslists = black.list.tld/<;192.168.1.2;a.domain
```

The DNS lookups that occur are:

```
2.1.168.192.black.list.tld
a.domain.black.list.tld
```

Once a DNS record has been found (that matches a specific IP return address, if specified – see section 43.33), no further lookups are done. If there is a temporary DNS error, the rest of the sublist of domains or IP addresses is tried. A temporary error for the whole dnslists item occurs only if no other DNS lookup in this sublist succeeds. In other words, a successful lookup for any of the items in the sublist overrides a temporary error for a previous item.

The ability to supply a list of items after the slash is in some sense just a syntactic convenience. These two examples have the same effect:

```
dnslists = black.list.tld/a.domain : black.list.tld/b.domain
dnslists = black.list.tld/a.domain::b.domain
```

However, when the data for the list is obtained from a lookup, the second form is usually much more convenient. Consider this example:

```
deny message = The mail servers for the domain \
               $sender_address_domain \
               are listed at $dnslist_domain ($dnslist_value); \
               see $dnslist_text.
dnslists = sbl.spamhaus.org/<|${lookup dnsdb {>|a=<|\
                               ${lookup dnsdb {>|mxh=\
                               $sender_address_domain} }} }
```

Note the use of >| in the dnsdb lookup to specify the separator for multiple DNS records. The inner dnsdb lookup produces a list of MX hosts and the outer dnsdb lookup finds the IP addresses for these hosts. The result of expanding the condition might be something like this:

```
dnslists = sbl.spamhaus.org/<|192.168.2.3|192.168.5.6|...
```

Thus, this example checks whether or not the IP addresses of the sender domain's mail servers are on the Spamhaus black list.

The key that was used for a successful DNS list lookup is put into the variable *\$dnslist\_matched* (see section 43.32).

### 43.31 Data returned by DNS lists

DNS lists are constructed using address records in the DNS. The original RBL just used the address 127.0.0.1 on the right hand side of each record, but the RBL+ list and some other lists use a number of values with different meanings. The values used on the RBL+ list are:

```
127.1.0.1 RBL
127.1.0.2 DUL
```

- 127.1.0.3 DUL and RBL
- 127.1.0.4 RSS
- 127.1.0.5 RSS and RBL
- 127.1.0.6 RSS and DUL
- 127.1.0.7 RSS and DUL and RBL

Section 43.33 below describes how you can distinguish between different values. Some DNS lists may return more than one address record; see section 43.35 for details of how they are checked.

### 43.32 Variables set from DNS lists

When an entry is found in a DNS list, the variable *\$dnslist\_domain* contains the name of the overall domain that matched (for example, *spamhaus.example*), *\$dnslist\_matched* contains the key within that domain (for example, *192.168.5.3*), and *\$dnslist\_value* contains the data from the DNS record. When the key is an IP address, it is not reversed in *\$dnslist\_matched* (though it is, of course, in the actual lookup). In simple cases, for example:

```
deny dnslists = spamhaus.example
```

the key is also available in another variable (in this case, *\$sender\_host\_address*). In more complicated cases, however, this is not true. For example, using a data lookup (as described in section 43.30) might generate a dnslists lookup like this:

```
deny dnslists = spamhaus.example/<|192.168.1.2|192.168.6.7|...
```

If this condition succeeds, the value in *\$dnslist\_matched* might be *192.168.6.7* (for example).

If more than one address record is returned by the DNS lookup, all the IP addresses are included in *\$dnslist\_value*, separated by commas and spaces. The variable *\$dnslist\_text* contains the contents of any associated TXT record. For lists such as RBL+ the TXT record for a merged entry is often not very meaningful. See section 43.36 for a way of obtaining more information.

You can use the DNS list variables in **message** or **log\_message** modifiers – although these appear before the condition in the ACL, they are not expanded until after it has failed. For example:

```
deny      hosts = !+local_networks
          message = $sender_host_address is listed \
                  at $dnslist_domain
          dnslists = rbl-plus.mail-abuse.example
```

### 43.33 Additional matching conditions for DNS lists

You can add an equals sign and an IP address after a **dnslists** domain name in order to restrict its action to DNS records with a matching right hand side. For example,

```
deny dnslists = rblplus.mail-abuse.org=127.0.0.2
```

rejects only those hosts that yield 127.0.0.2. Without this additional data, any address record is considered to be a match. For the moment, we assume that the DNS lookup returns just one record. Section 43.35 describes how multiple records are handled.

More than one IP address may be given for checking, using a comma as a separator. These are alternatives – if any one of them matches, the **dnslists** condition is true. For example:

```
deny dnslists = a.b.c=127.0.0.2,127.0.0.3
```

If you want to specify a constraining address list and also specify names or IP addresses to be looked up, the constraining address list must be specified first. For example:

```
deny dnslists = dsn.rfc-ignorant.org\
              =127.0.0.2/$sender_address_domain
```

If the character & is used instead of =, the comparison for each listed IP address is done by a bitwise “and” instead of by an equality test. In other words, the listed addresses are used as bit masks. The

comparison is true if all the bits in the mask are present in the address that is being tested. For example:

```
dnslists = a.b.c&0.0.0.3
```

matches if the address is *x.x.x.3*, *x.x.x.7*, *x.x.x.11*, etc. If you want to test whether one bit or another bit is present (as opposed to both being present), you must use multiple values. For example:

```
dnslists = a.b.c&0.0.0.1,0.0.0.2
```

matches if the final component of the address is an odd number or two times an odd number.

### 43.34 Negated DNS matching conditions

You can supply a negative list of IP addresses as part of a **dnslists** condition. Whereas

```
deny dnslists = a.b.c=127.0.0.2,127.0.0.3
```

means “deny if the host is in the black list at the domain *a.b.c* and the IP address yielded by the list is either 127.0.0.2 or 127.0.0.3”,

```
deny dnslists = a.b.c!=127.0.0.2,127.0.0.3
```

means “deny if the host is in the black list at the domain *a.b.c* and the IP address yielded by the list is not 127.0.0.2 and not 127.0.0.3”. In other words, the result of the test is inverted if an exclamation mark appears before the = (or the &) sign.

**Note:** This kind of negation is not the same as negation in a domain, host, or address list (which is why the syntax is different).

If you are using just one list, the negation syntax does not gain you much. The previous example is precisely equivalent to

```
deny dnslists = a.b.c
!dnslists = a.b.c=127.0.0.2,127.0.0.3
```

However, if you are using multiple lists, the negation syntax is clearer. Consider this example:

```
deny dnslists = sbl.spamhaus.org : \
list.dsbl.org : \
dnsbl.njabl.org!=127.0.0.3 : \
relays.ordb.org
```

Using only positive lists, this would have to be:

```
deny dnslists = sbl.spamhaus.org : \
list.dsbl.org
deny dnslists = dnsbl.njabl.org
!dnslists = dnsbl.njabl.org=127.0.0.3
deny dnslists = relays.ordb.org
```

which is less clear, and harder to maintain.

### 43.35 Handling multiple DNS records from a DNS list

A DNS lookup for a **dnslists** condition may return more than one DNS record, thereby providing more than one IP address. When an item in a **dnslists** list is followed by = or & and a list of IP addresses, in order to restrict the match to specific results from the DNS lookup, there are two ways in which the checking can be handled. For example, consider the condition:

```
dnslists = a.b.c=127.0.0.1
```

What happens if the DNS lookup for the incoming IP address yields both 127.0.0.1 and 127.0.0.2 by means of two separate DNS records? Is the condition true because at least one given value was found, or is it false because at least one of the found values was not listed? And how does this affect negated conditions? Both possibilities are provided for with the help of additional separators == and =&.

- If = or & is used, the condition is true if any one of the looked up IP addresses matches one of the listed addresses. For the example above, the condition is true because 127.0.0.1 matches.
- If == or =& is used, the condition is true only if every one of the looked up IP addresses matches one of the listed addresses. If the condition is changed to:

```
dnslists = a.b.c==127.0.0.1
```

and the DNS lookup yields both 127.0.0.1 and 127.0.0.2, the condition is false because 127.0.0.2 is not listed. You would need to have:

```
dnslists = a.b.c==127.0.0.1,127.0.0.2
```

for the condition to be true.

When ! is used to negate IP address matching, it inverts the result, giving the precise opposite of the behaviour above. Thus:

- If != or !& is used, the condition is true if none of the looked up IP addresses matches one of the listed addresses. Consider:

```
dnslists = a.b.c!&0.0.0.1
```

If the DNS lookup yields both 127.0.0.1 and 127.0.0.2, the condition is false because 127.0.0.1 matches.

- If != or !=& is used, the condition is true if there is at least one looked up IP address that does not match. Consider:

```
dnslists = a.b.c!=&0.0.0.1
```

If the DNS lookup yields both 127.0.0.1 and 127.0.0.2, the condition is true, because 127.0.0.2 does not match. You would need to have:

```
dnslists = a.b.c!=&0.0.0.1,0.0.0.2
```

for the condition to be false.

When the DNS lookup yields only a single IP address, there is no difference between = and == and between & and =&.

### 43.36 Detailed information from merged DNS lists

When the facility for restricting the matching IP values in a DNS list is used, the text from the TXT record that is set in *\$dnslist\_text* may not reflect the true reason for rejection. This happens when lists are merged and the IP address in the A record is used to distinguish them; unfortunately there is only one TXT record. One way round this is not to use merged lists, but that can be inefficient because it requires multiple DNS lookups where one would do in the vast majority of cases when the host of interest is not on any of the lists.

A less inefficient way of solving this problem is available. If two domain names, comma-separated, are given, the second is used first to do an initial check, making use of any IP value restrictions that are set. If there is a match, the first domain is used, without any IP value restrictions, to get the TXT record. As a byproduct of this, there is also a check that the IP being tested is indeed on the first list. The first domain is the one that is put in *\$dnslist\_domain*. For example:

```
reject message = \
    rejected because $sender_host_address is blacklisted \
    at $dnslist_domain\n$dnslist_text
dnslists = \
    sbl.spamhaus.org,sbl-xbl.spamhaus.org=127.0.0.2 : \
    dul.dnsbl.sorbs.net,dnsbl.sorbs.net=127.0.0.10
```

For the first blacklist item, this starts by doing a lookup in *sbl-xbl.spamhaus.org* and testing for a 127.0.0.2 return. If there is a match, it then looks in *sbl.spamhaus.org*, without checking the return value, and as long as something is found, it looks for the corresponding TXT record. If there is no

match in *sbl-xbl.spamhaus.org*, nothing more is done. The second blacklist item is processed similarly.

If you are interested in more than one merged list, the same list must be given several times, but because the results of the DNS lookups are cached, the DNS calls themselves are not repeated. For example:

```
reject dnslists = \  
    http.dnsbl.sorbs.net,dnsbl.sorbs.net=127.0.0.2 : \  
    socks.dnsbl.sorbs.net,dnsbl.sorbs.net=127.0.0.3 : \  
    misc.dnsbl.sorbs.net,dnsbl.sorbs.net=127.0.0.4 : \  
    dul.dnsbl.sorbs.net,dnsbl.sorbs.net=127.0.0.10
```

In this case there is one lookup in *dnsbl.sorbs.net*, and if none of the IP values matches (or if no record is found), this is the only lookup that is done. Only if there is a match is one of the more specific lists consulted.

### 43.37 DNS lists and IPv6

If Exim is asked to do a dnslist lookup for an IPv6 address, it inverts it nibble by nibble. For example, if the calling host's IP address is 3ffe:ffff:836f:0a00:000a:0800:200a:c031, Exim might look up

```
1.3.0.c.a.0.0.2.0.0.8.0.a.0.0.0.0.0.a.0.f.6.3.8.  
f.f.f.f.e.f.f.3.blackholes.mail-abuse.org
```

(split over two lines here to fit on the page). Unfortunately, some of the DNS lists contain wildcard records, intended for IPv4, that interact badly with IPv6. For example, the DNS entry

```
*.3.some.list.example.      A      127.0.0.1
```

is probably intended to put the entire 3.0.0.0/8 IPv4 network on the list. Unfortunately, it also matches the entire 3::4 IPv6 network.

You can exclude IPv6 addresses from DNS lookups by making use of a suitable **condition** condition, as in this example:

```
deny    condition = ${if isip4{$sender_host_address}}  
        dnslists  = some.list.example
```

If an explicit key is being used for a DNS lookup and it may be an IPv6 address you should specify alternate list separators for both the outer (DNS list name) list and inner (lookup keys) list:

```
dnslists = <; dnsbl.example.com/<|$acl_m_addrslst
```

### 43.38 Rate limiting incoming messages

The **ratelimit** ACL condition can be used to measure and control the rate at which clients can send email. This is more powerful than the **smtp\_ratelimit\_\*** options, because those options control the rate of commands in a single SMTP session only, whereas the **ratelimit** condition works across all connections (concurrent and sequential) from the same client host. The syntax of the **ratelimit** condition is:

```
ratelimit = <m> / <p> / <options> / <key>
```

If the average client sending rate is less than *m* messages per time period *p* then the condition is false; otherwise it is true.

As a side-effect, the **ratelimit** condition sets the expansion variable *\$sender\_rate* to the client's computed rate, *\$sender\_rate\_limit* to the configured value of *m*, and *\$sender\_rate\_period* to the configured value of *p*.

The parameter *p* is the smoothing time constant, in the form of an Exim time interval, for example, 8h for eight hours. A larger time constant means that it takes Exim longer to forget a client's past behaviour. The parameter *m* is the maximum number of messages that a client is permitted to send in each time interval. It also specifies the number of messages permitted in a fast burst. By increasing



both *m* and *p* but keeping *m/p* constant, you can allow a client to send more messages in a burst without changing its long-term sending rate limit. Conversely, if *m* and *p* are both small, messages must be sent at an even rate.

There is a script in *util/ratelimit.pl* which extracts sending rates from log files, to assist with choosing appropriate settings for *m* and *p* when deploying the **ratelimit** ACL condition. The script prints usage instructions when it is run with no arguments.

The key is used to look up the data for calculating the client's average sending rate. This data is stored in Exim's spool directory, alongside the retry and other hints databases. The default key is *\$sender\_host\_address*, which means Exim computes the sending rate of each client host IP address. By changing the key you can change how Exim identifies clients for the purpose of ratelimiting. For example, to limit the sending rate of each authenticated user, independent of the computer they are sending from, set the key to *\$authenticated\_id*. You must ensure that the lookup key is meaningful; for example, *\$authenticated\_id* is only meaningful if the client has authenticated (which you can check with the **authenticated** ACL condition).

The lookup key does not have to identify clients: If you want to limit the rate at which a recipient receives messages, you can use the key *\$local\_part@\$domain* with the **per\_rcpt** option (see below) in a RCPT ACL.

Each **ratelimit** condition can have up to four options. A **per\_\*** option specifies what Exim measures the rate of, for example messages or recipients or bytes. You can adjust the measurement using the **unique=** and/or **count=** options. You can also control when Exim updates the recorded rate using a **strict**, **leaky**, or **readonly** option. The options are separated by a slash, like the other parameters. They may appear in any order.

Internally, Exim appends the smoothing constant *p* onto the lookup key with any options that alter the meaning of the stored data. The limit *m* is not stored, so you can alter the configured maximum rate and Exim will still remember clients' past behaviour. If you change the **per\_\*** mode or add or remove the **unique=** option, the lookup key changes so Exim will forget past behaviour. The lookup key is not affected by changes to the update mode and the **count=** option.

### 43.39 Ratelimit options for what is being measured

The **per\_conn** option limits the client's connection rate. It is not normally used in the **acl\_not\_smtp**, **acl\_not\_smtp\_mime**, or **acl\_not\_smtp\_start** ACLs.

The **per\_mail** option limits the client's rate of sending messages. This is the default if none of the **per\_\*** options is specified. It can be used in **acl\_smtp\_mail**, **acl\_smtp\_rcpt**, **acl\_smtp\_predata**, **acl\_smtp\_mime**, **acl\_smtp\_data**, or **acl\_not\_smtp**.

The **per\_byte** option limits the sender's email bandwidth. It can be used in the same ACLs as the **per\_mail** option, though it is best to use this option in the **acl\_smtp\_mime**, **acl\_smtp\_data** or **acl\_not\_smtp** ACLs; if it is used in an earlier ACL, Exim relies on the SIZE parameter given by the client in its MAIL command, which may be inaccurate or completely missing. You can follow the limit *m* in the configuration with K, M, or G to specify limits in kilobytes, megabytes, or gigabytes, respectively.

The **per\_rcpt** option causes Exim to limit the rate at which recipients are accepted. It can be used in the **acl\_smtp\_rcpt**, **acl\_smtp\_predata**, **acl\_smtp\_mime**, **acl\_smtp\_data**, or **acl\_smtp\_rcpt** ACLs. In **acl\_smtp\_rcpt** the rate is updated one recipient at a time; in the other ACLs the rate is updated with the total (accepted) recipient count in one go. Note that in either case the rate limiting engine will see a message with many recipients as a large high-speed burst.

The **per\_addr** option is like the **per\_rcpt** option, except it counts the number of different recipients that the client has sent messages to in the last time period. That is, if the client repeatedly sends messages to the same recipient, its measured rate is not increased. This option can only be used in **acl\_smtp\_rcpt**.

The **per\_cmd** option causes Exim to recompute the rate every time the condition is processed. This can be used to limit the rate of any SMTP command. If it is used in multiple ACLs it can limit the aggregate rate of multiple different commands.

The **count=** option can be used to alter how much Exim adds to the client's measured rate. For example, the **per\_byte** option is equivalent to `per_mail/count=$message_size`. If there is no **count=** option, Exim increases the measured rate by one (except for the **per\_rcpt** option in ACLs other than **acl\_smtp\_rcpt**). The count does not have to be an integer.

The **unique=** option is described in section 43.42 below.

## 43.40 Ratelimit update modes

You can specify one of three options with the **ratelimit** condition to control when its database is updated. This section describes the **readonly** mode, and the next section describes the **strict** and **leaky** modes.

If the **ratelimit** condition is used in **readonly** mode, Exim looks up a previously-computed rate to check against the limit.

For example, you can test the client's sending rate and deny it access (when it is too fast) in the connect ACL. If the client passes this check then it can go on to send a message, in which case its recorded rate will be updated in the MAIL ACL. Subsequent connections from the same client will check this new rate.

```
acl_check_connect:
  deny ratelimit = 100 / 5m / readonly
    log_message = RATE CHECK: $sender_rate/$sender_rate_period \
                  (max $sender_rate_limit)
# ...
acl_check_mail:
  warn ratelimit = 100 / 5m / strict
    log_message = RATE UPDATE: $sender_rate/$sender_rate_period \
                  (max $sender_rate_limit)
```

If Exim encounters multiple **ratelimit** conditions with the same key when processing a message then it may increase the client's measured rate more than it should. For example, this will happen if you check the **per\_rcpt** option in both **acl\_smtp\_rcpt** and **acl\_smtp\_data**. However it's OK to check the same **ratelimit** condition multiple times in the same ACL. You can avoid any multiple update problems by using the **readonly** option on later ratelimit checks.

The **per\_\*** options described above do not make sense in some ACLs. If you use a **per\_\*** option in an ACL where it is not normally permitted then the update mode defaults to **readonly** and you cannot specify the **strict** or **leaky** modes. In other ACLs the default update mode is **leaky** (see the next section) so you must specify the **readonly** option explicitly.

## 43.41 Ratelimit options for handling fast clients

If a client's average rate is greater than the maximum, the rate limiting engine can react in two possible ways, depending on the presence of the **strict** or **leaky** update modes. This is independent of the other counter-measures (such as rejecting the message) that may be specified by the rest of the ACL.

The **leaky** (default) option means that the client's recorded rate is not updated if it is above the limit. The effect of this is that Exim measures the client's average rate of successfully sent email, which cannot be greater than the maximum allowed. If the client is over the limit it may suffer some counter-measures (as specified in the ACL), but it will still be able to send email at the configured maximum rate, whatever the rate of its attempts. This is generally the better choice if you have clients that retry automatically. For example, it does not prevent a sender with an over-aggressive retry rate from getting any email through.

The **strict** option means that the client's recorded rate is always updated. The effect of this is that Exim measures the client's average rate of attempts to send email, which can be much higher than the maximum it is actually allowed. If the client is over the limit it may be subjected to counter-measures by the ACL. It must slow down and allow sufficient time to pass that its computed rate falls below the

maximum before it can send email again. The time (the number of smoothing periods) it must wait and not attempt to send mail can be calculated with this formula:

$$\ln(\text{peakrate}/\text{maxrate})$$

## 43.42 Limiting the rate of different events

The **ratelimit unique=** option controls a mechanism for counting the rate of different events. For example, the **per\_addr** option uses this mechanism to count the number of different recipients that the client has sent messages to in the last time period; it is equivalent to `per_rcpt/unique=$local_part@$domain`. You could use this feature to measure the rate that a client uses different sender addresses with the options `per_mail/unique=$sender_address`.

For each **ratelimit** key Exim stores the set of **unique=** values that it has seen for that key. The whole set is thrown away when it is older than the rate smoothing period *p*, so each different event is counted at most once per period. In the **leaky** update mode, an event that causes the client to go over the limit is not added to the set, in the same way that the client's recorded rate is not updated in the same situation.

When you combine the **unique=** and **readonly** options, the specific **unique=** value is ignored, and Exim just retrieves the client's stored rate.

The **unique=** mechanism needs more space in the ratelimit database than the other **ratelimit** options in order to store the event set. The number of unique values is potentially as large as the rate limit, so the extra space required increases with larger limits.

The uniqueification is not perfect: there is a small probability that Exim will think a new event has happened before. If the sender's rate is less than the limit, Exim should be more than 99.9% correct. However in **strict** mode the measured rate can go above the limit, in which case Exim may undercount events by a significant margin. Fortunately, if the rate is high enough (2.7 times the limit) that the false positive rate goes above 9%, then Exim will throw away the over-full event set before the measured rate falls below the limit. Therefore the only harm should be that exceptionally high sending rates are logged incorrectly; any countermeasures you configure will be as effective as intended.

## 43.43 Using rate limiting

Exim's other ACL facilities are used to define what counter-measures are taken when the rate limit is exceeded. This might be anything from logging a warning (for example, while measuring existing sending rates in order to define policy), through time delays to slow down fast senders, up to rejecting the message. For example:

```
# Log all senders' rates
warn ratelimit = 0 / 1h / strict
    log_message = Sender rate $sender_rate / $sender_rate_period

# Slow down fast senders; note the need to truncate $sender_rate
# at the decimal point.
warn ratelimit = 100 / 1h / per_rcpt / strict
    delay      = ${eval: ${sg{$sender_rate}}{[.]*}} - \
                $sender_rate_limit }s

# Keep authenticated users under control
deny authenticated = *
    ratelimit = 100 / 1d / strict / $authenticated_id

# System-wide rate limit
defer message = Sorry, too busy. Try again later.
    ratelimit = 10 / 1s / $primary_hostname

# Restrict incoming rate from each host, with a default
```

```
# set using a macro and special cases looked up in a table.
defer message = Sender rate exceeds $sender_rate_limit \
    messages per $sender_rate_period
ratelimit = ${lookup {$sender_host_address} \
    cdb {DB/ratelimits.cdb} \
    {$value} {RATELIMIT} }
```

**Warning:** If you have a busy server with a lot of **ratelimit** tests, especially with the **per\_rcpt** option, you may suffer from a performance bottleneck caused by locking on the ratelimit hints database. Apart from making your ACLs less complicated, you can reduce the problem by using a RAM disk for Exim's hints directory (usually */var/spool/exim/db/*). However this means that Exim will lose its hints data after a reboot (including retry hints, the callout cache, and ratelimit data).

## 43.44 Address verification

Several of the **verify** conditions described in section 43.26 cause addresses to be verified. Section 43.48 discusses the reporting of sender verification failures. The verification conditions can be followed by options that modify the verification process. The options are separated from the keyword and from each other by slashes, and some of them contain parameters. For example:

```
verify = sender/callout
verify = recipient/defer_ok/callout=10s,defer_ok
```

The first stage of address verification, which always happens, is to run the address through the routers, in “verify mode”. Routers can detect the difference between verification and routing for delivery, and their actions can be varied by a number of generic options such as **verify** and **verify\_only** (see chapter 15). If routing fails, verification fails. The available options are as follows:

- If the **callout** option is specified, successful routing to one or more remote hosts is followed by a “callout” to those hosts as an additional check. Callouts and their sub-options are discussed in the next section.
- If there is a defer error while doing verification routing, the ACL normally returns “defer”. However, if you include **defer\_ok** in the options, the condition is forced to be true instead. Note that this is a main verification option as well as a suboption for callouts.
- The **no\_details** option is covered in section 43.48, which discusses the reporting of sender address verification failures.
- The **success\_on\_redirect** option causes verification always to succeed immediately after a successful redirection. By default, if a redirection generates just one address, that address is also verified. See further discussion in section 43.49.

After an address verification failure, *\$acl\_verify\_message* contains the error message that is associated with the failure. It can be preserved by coding like this:

```
warn !verify = sender
    set acl_m0 = $acl_verify_message
```

If you are writing your own custom rejection message or log message when denying access, you can use this variable to include information about the verification failure.

In addition, *\$sender\_verify\_failure* or *\$recipient\_verify\_failure* (as appropriate) contains one of the following words:

- **qualify**: The address was unqualified (no domain), and the message was neither local nor came from an exempted host.
- **route**: Routing failed.
- **mail**: Routing succeeded, and a callout was attempted; rejection occurred at or before the MAIL command (that is, on initial connection, HELO, or MAIL).
- **recipient**: The RCPT command in a callout was rejected.
- **postmaster**: The postmaster check in a callout was rejected.

The main use of these variables is expected to be to distinguish between rejections of MAIL and rejections of RCPT in callouts.

### 43.45 Callout verification

For non-local addresses, routing verifies the domain, but is unable to do any checking of the local part. There are situations where some means of verifying the local part is desirable. One way this can be done is to make an SMTP *callback* to a delivery host for the sender address or a *callforward* to a subsequent host for a recipient address, to see if the host accepts the address. We use the term *callout* to cover both cases. Note that for a sender address, the callback is not to the client host that is trying to deliver the message, but to one of the hosts that accepts incoming mail for the sender's domain.

Exim does not do callouts by default. If you want them to happen, you must request them by setting appropriate options on the **verify** condition, as described below. This facility should be used with care, because it can add a lot of resource usage to the cost of verifying an address. However, Exim does cache the results of callouts, which helps to reduce the cost. Details of caching are in section 43.47.

Recipient callouts are usually used only between hosts that are controlled by the same administration. For example, a corporate gateway host could use callouts to check for valid recipients on an internal mailserver. A successful callout does not guarantee that a real delivery to the address would succeed; on the other hand, a failing callout does guarantee that a delivery would fail.

If the **callout** option is present on a condition that verifies an address, a second stage of verification occurs if the address is successfully routed to one or more remote hosts. The usual case is routing by a *dnslookup* or a *manualroute* router, where the router specifies the hosts. However, if a router that does not set up hosts routes to an *smtp* transport with a **hosts** setting, the transport's hosts are used. If an *smtp* transport has **hosts\_override** set, its hosts are always used, whether or not the router supplies a host list. Callouts are only supported on *smtp* transports.

The port that is used is taken from the transport, if it is specified and is a remote transport. (For routers that do verification only, no transport need be specified.) Otherwise, the default SMTP port is used. If a remote transport specifies an outgoing interface, this is used; otherwise the interface is not specified. Likewise, the text that is used for the HELO command is taken from the transport's **helo\_data** option; if there is no transport, the value of *\$smtp\_active\_hostname* is used.

For a sender callout check, Exim makes SMTP connections to the remote hosts, to test whether a bounce message could be delivered to the sender address. The following SMTP commands are sent:

```
HELO <local host name>
MAIL FROM:<>
RCPT TO:<the address to be tested>
QUIT
```

LHLO is used instead of HELO if the transport's **protocol** option is set to "lmtpl".

The callout may use EHLO, AUTH and/or STARTTLS given appropriate option settings.

A recipient callout check is similar. By default, it also uses an empty address for the sender. This default is chosen because most hosts do not make use of the sender address when verifying a recipient. Using the same address means that a single cache entry can be used for each recipient. Some sites, however, do make use of the sender address when verifying. These are catered for by the **use\_sender** and **use\_postmaster** options, described in the next section.

If the response to the RCPT command is a 2xx code, the verification succeeds. If it is 5xx, the verification fails. For any other condition, Exim tries the next host, if any. If there is a problem with all the remote hosts, the ACL yields "defer", unless the **defer\_ok** parameter of the **callout** option is given, in which case the condition is forced to succeed.

A callout may take a little time. For this reason, Exim normally flushes SMTP output before performing a callout in an ACL, to avoid unexpected timeouts in clients when the SMTP PIPELINING extension is in use. The flushing can be disabled by using a **control** modifier to set **no\_callout\_flush**.

## 43.46 Additional parameters for callouts

The **callout** option can be followed by an equals sign and a number of optional parameters, separated by commas. For example:

```
verify = recipient/callout=10s,defer_ok
```

The old syntax, which had **callout\_defer\_ok** and **check\_postmaster** as separate verify options, is retained for backwards compatibility, but is now deprecated. The additional parameters for **callout** are as follows:

**<a time interval>**

This specifies the timeout that applies for the callout attempt to each host. For example:

```
verify = sender/callout=5s
```

The default is 30 seconds. The timeout is used for each response from the remote host. It is also used for the initial connection, unless overridden by the **connect** parameter.

**connect = <time interval>**

This parameter makes it possible to set a different (usually smaller) timeout for making the SMTP connection. For example:

```
verify = sender/callout=5s,connect=1s
```

If not specified, this timeout defaults to the general timeout value.

**defer\_ok**

When this parameter is present, failure to contact any host, or any other kind of temporary error, is treated as success by the ACL. However, the cache is not updated in this circumstance.

**fullpostmaster**

This operates like the **postmaster** option (see below), but if the check for *postmaster@domain* fails, it tries just *postmaster*, without a domain, in accordance with the specification in RFC 2821. The RFC states that the unqualified address *postmaster* should be accepted.

**mailfrom = <email address>**

When verifying addresses in header lines using the **header\_sender** verification option, Exim behaves by default as if the addresses are envelope sender addresses from a message. Callout verification therefore tests to see whether a bounce message could be delivered, by using an empty address in the MAIL command. However, it is arguable that these addresses might never be used as envelope senders, and could therefore justifiably reject bounce messages (empty senders). The **mailfrom** callout parameter allows you to specify what address to use in the MAIL command. For example:

```
require verify = header_sender/callout=mailfrom=abcd@x.y.z
```

This parameter is available only for the **header\_sender** verification option.

**maxwait = <time interval>**

This parameter sets an overall timeout for performing a callout verification. For example:

```
verify = sender/callout=5s,maxwait=30s
```

This timeout defaults to four times the callout timeout for individual SMTP commands. The overall timeout applies when there is more than one host that can be tried. The timeout is checked before trying the next host. This prevents very long delays if there are a large number of hosts and all are timing out (for example, when network connections are timing out).

**no\_cache**

When this parameter is given, the callout cache is neither read nor updated.

**postmaster**

When this parameter is set, a successful callout check is followed by a similar check for the local part *postmaster* at the same domain. If this address is rejected, the callout fails (but see **fullpostmaster** above). The result of the postmaster check is recorded in a cache record; if it is a

failure, this is used to fail subsequent callouts for the domain without a connection being made, until the cache record expires.

#### **postmaster\_mailfrom** = <email address>

The postmaster check uses an empty sender in the MAIL command by default. You can use this parameter to do a postmaster check using a different address. For example:

```
require verify = sender/callout=postmaster_mailfrom=abc@x.y.z
```

If both **postmaster** and **postmaster\_mailfrom** are present, the rightmost one overrides. The **postmaster** parameter is equivalent to this example:

```
require verify = sender/callout=postmaster_mailfrom=
```

**Warning:** The caching arrangements for postmaster checking do not take account of the sender address. It is assumed that either the empty address or a fixed non-empty address will be used. All that Exim remembers is that the postmaster check for the domain succeeded or failed.

#### **random**

When this parameter is set, before doing the normal callout check, Exim does a check for a “random” local part at the same domain. The local part is not really random – it is defined by the expansion of the option **callout\_random\_local\_part**, which defaults to

```
$primary_hostname-$tod_epoch-testing
```

The idea here is to try to determine whether the remote host accepts all local parts without checking. If it does, there is no point in doing callouts for specific local parts. If the “random” check succeeds, the result is saved in a cache record, and used to force the current and subsequent callout checks to succeed without a connection being made, until the cache record expires.

#### **use\_postmaster**

This parameter applies to recipient callouts only. For example:

```
deny !verify = recipient/callout=use_postmaster
```

It causes a non-empty postmaster address to be used in the MAIL command when performing the callout for the recipient, and also for a “random” check if that is configured. The local part of the address is *postmaster* and the domain is the contents of *\$qualify\_domain*.

#### **use\_sender**

This option applies to recipient callouts only. For example:

```
require verify = recipient/callout=use_sender
```

It causes the message’s actual sender address to be used in the MAIL command when performing the callout, instead of an empty address. There is no need to use this option unless you know that the called hosts make use of the sender when checking recipients. If used indiscriminately, it reduces the usefulness of callout caching.

#### **hold**

This option applies to recipient callouts only. For example:

```
require verify = recipient/callout=use_sender,hold
```

It causes the connection to be held open and used for any further recipients and for eventual delivery (should that be done quickly). Doing this saves on TCP and SMTP startup costs, and TLS costs also when that is used for the connections. The advantage is only gained if there are no callout cache hits (which could be enforced by the *no\_cache* option), if the *use\_sender* option is used, if neither the *random* nor the *use\_postmaster* option is used, and if no other callouts intervene.

If you use any of the parameters that set a non-empty sender for the MAIL command (**mailfrom**, **postmaster\_mailfrom**, **use\_postmaster**, or **use\_sender**), you should think about possible loops. Recipient checking is usually done between two hosts that are under the same management, and the host that receives the callouts is not normally configured to do callouts itself. Therefore, it is normally safe to use **use\_postmaster** or **use\_sender** in these circumstances.

However, if you use a non-empty sender address for a callout to an arbitrary host, there is the likelihood that the remote host will itself initiate a callout check back to your host. As it is checking what appears to be a message sender, it is likely to use an empty address in MAIL, thus avoiding a callout loop. However, to be on the safe side it would be best to set up your own ACLs so that they do not do sender verification checks when the recipient is the address you use for header sender or postmaster callout checking.

Another issue to think about when using non-empty senders for callouts is caching. When you set **mailfrom** or **use\_sender**, the cache record is keyed by the sender/recipient combination; thus, for any given recipient, many more actual callouts are performed than when an empty sender or postmaster is used.

### 43.47 Callout caching

Exim caches the results of callouts in order to reduce the amount of resources used, unless you specify the **no\_cache** parameter with the **callout** option. A hints database called “callout” is used for the cache. Two different record types are used: one records the result of a callout check for a specific address, and the other records information that applies to the entire domain (for example, that it accepts the local part *postmaster*).

When an original callout fails, a detailed SMTP error message is given about the failure. However, for subsequent failures use the cache data, this message is not available.

The expiry times for negative and positive address cache records are independent, and can be set by the global options **callout\_negative\_expire** (default 2h) and **callout\_positive\_expire** (default 24h), respectively.

If a host gives a negative response to an SMTP connection, or rejects any commands up to and including

```
MAIL FROM:<>
```

(but not including the MAIL command with a non-empty address), any callout attempt is bound to fail. Exim remembers such failures in a domain cache record, which it uses to fail callouts for the domain without making new connections, until the domain record times out. There are two separate expiry times for domain cache records: **callout\_domain\_negative\_expire** (default 3h) and **callout\_domain\_positive\_expire** (default 7d).

Domain records expire when the negative expiry time is reached if callouts cannot be made for the domain, or if the postmaster check failed. Otherwise, they expire when the positive expiry time is reached. This ensures that, for example, a host that stops accepting “random” local parts will eventually be noticed.

The callout caching mechanism is based on the domain of the address that is being tested. If the domain routes to several hosts, it is assumed that their behaviour will be the same.

### 43.48 Sender address verification reporting

See section 43.44 for a general discussion of verification. When sender verification fails in an ACL, the details of the failure are given as additional output lines before the 550 response to the relevant SMTP command (RCPT or DATA). For example, if sender callout is in use, you might see:

```
MAIL FROM:<xyz@abc.example>
250 OK
RCPT TO:<pqr@def.example>
550-Verification failed for <xyz@abc.example>
550-Called: 192.168.34.43
550-Sent: RCPT TO:<xyz@abc.example>
550-Response: 550 Unknown local part xyz in <xyz@abc.example>
550 Sender verification failed
```

If more than one RCPT command fails in the same way, the details are given only for the first of them. However, some administrators do not want to send out this much information. You can suppress



the details by adding `/no_details` to the ACL statement that requests sender verification. For example:

```
verify = sender/no_details
```

### 43.49 Redirection while verifying

A dilemma arises when a local address is redirected by aliasing or forwarding during verification: should the generated addresses themselves be verified, or should the successful expansion of the original address be enough to verify it? By default, Exim takes the following pragmatic approach:

- When an incoming address is redirected to just one child address, verification continues with the child address, and if that fails to verify, the original verification also fails.
- When an incoming address is redirected to more than one child address, verification does not continue. A success result is returned.

This seems the most reasonable behaviour for the common use of aliasing as a way of redirecting different local parts to the same mailbox. It means, for example, that a pair of alias entries of the form

```
A.Wol:    aw123
aw123:    :fail: Gone away, no forwarding address
```

work as expected, with both local parts causing verification failure. When a redirection generates more than one address, the behaviour is more like a mailing list, where the existence of the alias itself is sufficient for verification to succeed.

It is possible, however, to change the default behaviour so that all successful redirections count as successful verifications, however many new addresses are generated. This is specified by the **success\_on\_redirect** verification option. For example:

```
require verify = recipient/success_on_redirect/callout=10s
```

In this example, verification succeeds if a router generates a new address, and the callout does not occur, because no address was routed to a remote host.

When verification is being tested via the **-bv** option, the treatment of redirections is as just described, unless the **-v** or any debugging option is also specified. In that case, full verification is done for every generated address and a report is output for each of them.

### 43.50 Client SMTP authorization (CSA)

Client SMTP Authorization is a system that allows a site to advertise which machines are and are not permitted to send email. This is done by placing special SRV records in the DNS; these are looked up using the client's HELO domain. At the time of writing, CSA is still an Internet Draft. Client SMTP Authorization checks in Exim are performed by the ACL condition:

```
verify = csa
```

This fails if the client is not authorized. If there is a DNS problem, or if no valid CSA SRV record is found, or if the client is authorized, the condition succeeds. These three cases can be distinguished using the expansion variable `$csa_status`, which can take one of the values “fail”, “defer”, “unknown”, or “ok”. The condition does not itself defer because that would be likely to cause problems for legitimate email.

The error messages produced by the CSA code include slightly more detail. If `$csa_status` is “defer”, this may be because of problems looking up the CSA SRV record, or problems looking up the CSA target address record. There are four reasons for `$csa_status` being “fail”:

- The client's host name is explicitly not authorized.
- The client's IP address does not match any of the CSA target IP addresses.
- The client's host name is authorized but it has no valid target IP addresses (for example, the target's addresses are IPv6 and the client is using IPv4).

- The client's host name has no CSA SRV record but a parent domain has asserted that all subdomains must be explicitly authorized.

The **csa** verification condition can take an argument which is the domain to use for the DNS query. The default is:

```
verify = csa/$sender_helo_name
```

This implementation includes an extension to CSA. If the query domain is an address literal such as [192.0.2.95], or if it is a bare IP address, Exim searches for CSA SRV records in the reverse DNS as if the HELO domain was (for example) *95.2.0.192.in-addr.arpa*. Therefore it is meaningful to say:

```
verify = csa/$sender_host_address
```

In fact, this is the check that Exim performs if the client does not say HELO. This extension can be turned off by setting the main configuration option **dns\_csa\_use\_reverse** to be false.

If a CSA SRV record is not found for the domain itself, a search is performed through its parent domains for a record which might be making assertions about subdomains. The maximum depth of this search is limited using the main configuration option **dns\_csa\_search\_limit**, which is 5 by default. Exim does not look for CSA SRV records in a top level domain, so the default settings handle HELO domains as long as seven (*hostname.five.four.three.two.one.com*). This encompasses the vast majority of legitimate HELO domains.

The *dnsdb* lookup also has support for CSA. Although *dnsdb* also supports direct SRV lookups, this is not sufficient because of the extra parent domain search behaviour of CSA, and (as with PTR lookups) *dnsdb* also turns IP addresses into lookups in the reverse DNS space. The result of a successful lookup such as:

```
${lookup dnsdb {csa=$sender_helo_name}}
```

has two space-separated fields: an authorization code and a target host name. The authorization code can be “Y” for yes, “N” for no, “X” for explicit authorization required but absent, or “?” for unknown.

## 43.51 Bounce address tag validation

Bounce address tag validation (BATV) is a scheme whereby the envelope senders of outgoing messages have a cryptographic, timestamped “tag” added to them. Genuine incoming bounce messages should therefore always be addressed to recipients that have a valid tag. This scheme is a way of detecting unwanted bounce messages caused by sender address forgeries (often called “collateral spam”), because the recipients of such messages do not include valid tags.

There are two expansion items to help with the implementation of the BATV “prvs” (private signature) scheme in an Exim configuration. This scheme signs the original envelope sender address by using a simple key to add a hash of the address and some time-based randomizing information. The **prvs** expansion item creates a signed address, and the **prvscheck** expansion item checks one. The syntax of these expansion items is described in section 11.5. The validity period on signed addresses is seven days.

As an example, suppose the secret per-address keys are stored in an MySQL database. A query to look up the key for an address could be defined as a macro like this:

```
PRVSCHECK_SQL = ${lookup mysql{SELECT secret FROM batv_prvs \
    WHERE sender='${quote_mysql:$prvscheck_address}'\
    }}{$value}}
```

Suppose also that the senders who make use of BATV are defined by an address list called **batv\_senders**. Then, in the ACL for RCPT commands, you could use this:

```
# Bounces: drop unsigned addresses for BATV senders
deny message = This address does not send an unsigned reverse path
  senders = :
  recipients = +batv_senders
```

```
# Bounces: In case of prvs-signed address, check signature.
deny message = Invalid reverse path signature.
senders = :
condition = ${prvscheck ${local_part}${domain}}\
            {PRVSCHECK_SQL}{1}}
!condition = $prvscheck_result
```

The first statement rejects recipients for bounce messages that are addressed to plain BATV sender addresses, because it is known that BATV senders do not send out messages with plain sender addresses. The second statement rejects recipients that are prvs-signed, but with invalid signatures (either because the key is wrong, or the signature has timed out).

A non-prvs-signed address is not rejected by the second statement, because the **prvscheck** expansion yields an empty string if its first argument is not a prvs-signed address, thus causing the **condition** to be false. If the first argument is a syntactically valid prvs-signed address, the yield is the third string (in this case “1”), whether or not the cryptographic and timeout checks succeed. The *\$prvscheck\_result* variable contains the result of the checks (empty for failure, “1” for success).

There is one more issue you must consider when implementing prvs-signing: you have to ensure that the routers accept prvs-signed addresses and deliver them correctly. The easiest way to handle this is to use a *redirect* router to remove the signature with a configuration along these lines:

```
batv_redirect:
    driver = redirect
    data = ${prvscheck ${local_part}${domain}}{PRVSCHECK_SQL}}
```

This works because, if the third argument of **prvscheck** is empty, the result of the expansion of a prvs-signed address is the decoded value of the original address. This router should probably be the first of your routers that handles local addresses.

To create BATV-signed addresses in the first place, a transport of this form can be used:

```
external_smtp_batv:
    driver = smtp
    return_path = ${prvs ${return_path} \
                        ${lookup mysql{SELECT \
                        secret FROM batv_prvs WHERE \
                        sender='${quote_mysql:$sender_address}'} \
                        {$value}fail}}}
```

If no key can be found for the existing return path, no signing takes place.

## 43.52 Using an ACL to control relaying

An MTA is said to *relay* a message if it receives it from some host and delivers it directly to another host as a result of a remote address contained within it. Redirecting a local address via an alias or forward file and then passing the message on to another host is not relaying, but a redirection as a result of the “percent hack” is.

Two kinds of relaying exist, which are termed “incoming” and “outgoing”. A host which is acting as a gateway or an MX backup is concerned with incoming relaying from arbitrary hosts to a specific set of domains. On the other hand, a host which is acting as a smart host for a number of clients is concerned with outgoing relaying from those clients to the Internet at large. Often the same host is fulfilling both functions, but in principle these two kinds of relaying are entirely independent. What is not wanted is the transmission of mail from arbitrary remote hosts through your system to arbitrary domains.

You can implement relay control by means of suitable statements in the ACL that runs for each RCPT command. For convenience, it is often easiest to use Exim’s named list facility to define the domains and hosts involved. For example, suppose you want to do the following:

- Deliver a number of domains to mailboxes on the local host (or process them locally in some other way). Let’s say these are *my.dom1.example* and *my.dom2.example*.

- Relay mail for a number of other domains for which you are the secondary MX. These might be *friend1.example* and *friend2.example*.
- Relay mail from the hosts on your local LAN, to whatever domains are involved. Suppose your LAN is 192.168.45.0/24.

In the main part of the configuration, you put the following definitions:

```
domainlist local_domains      = my.dom1.example : my.dom2.example
domainlist relay_to_domains   = friend1.example : friend2.example
hostlist   relay_from_hosts   = 192.168.45.0/24
```

Now you can use these definitions in the ACL that is run for every RCPT command:

```
acl_check_rcpt:
    accept domains = +local_domains : +relay_to_domains
    accept hosts   = +relay_from_hosts
```

The first statement accepts any RCPT command that contains an address in the local or relay domains. For any other domain, control passes to the second statement, which accepts the command only if it comes from one of the relay hosts. In practice, you will probably want to make your ACL more sophisticated than this, for example, by including sender and recipient verification. The default configuration includes a more comprehensive example, which is described in chapter 7.

### 43.53 Checking a relay configuration

You can check the relay characteristics of your configuration in the same way that you can test any ACL behaviour for an incoming SMTP connection, by using the **-bh** option to run a fake SMTP session with which you interact.

## 44. Content scanning at ACL time

The extension of Exim to include content scanning at ACL time, formerly known as “exiscan”, was originally implemented as a patch by Tom Kistner. The code was integrated into the main source for Exim release 4.50, and Tom continues to maintain it. Most of the wording of this chapter is taken from Tom’s specification.

It is also possible to scan the content of messages at other times. The *local\_scan()* function (see chapter 45) allows for content scanning after all the ACLs have run. A transport filter can be used to scan messages at delivery time (see the **transport\_filter** option, described in chapter 24).

If you want to include the ACL-time content-scanning features when you compile Exim, you need to arrange for **WITH\_CONTENT\_SCAN** to be defined in your *Local/Makefile*. When you do that, the Exim binary is built with:

- Two additional ACLs (**acl\_smtp\_mime** and **acl\_not\_smtp\_mime**) that are run for all MIME parts for SMTP and non-SMTP messages, respectively.
- Additional ACL conditions and modifiers: **decode**, **malware**, **mime\_regex**, **regex**, and **spam**. These can be used in the ACL that is run at the end of message reception (the **acl\_smtp\_data** ACL).
- An additional control feature (“no\_mbox\_unspool”) that saves spooled copies of messages, or parts of messages, for debugging purposes.
- Additional expansion variables that are set in the new ACL and by the new conditions.
- Two new main configuration options: **av\_scanner** and **spamd\_address**.

Content-scanning is continually evolving, and new features are still being added. While such features are still unstable and liable to incompatible changes, they are made available in Exim by setting options whose names begin **EXPERIMENTAL\_** in *Local/Makefile*. Such features are not documented in this manual. You can find out about them by reading the file called *doc/experimental.txt*.

All the content-scanning facilities work on a MBOX copy of the message that is temporarily created in a file called:

```
<spool_directory>/scan/<message_id>/<message_id>.eml
```

The *.eml* extension is a friendly hint to virus scanners that they can expect an MBOX-like structure inside that file. The file is created when the first content scanning facility is called. Subsequent calls to content scanning conditions open the same file again. The directory is recursively removed when the **acl\_smtp\_data** ACL has finished running, unless

```
control = no_mbox_unspool
```

has been encountered. When the MIME ACL decodes files, they are put into the same directory by default.

### 44.1 Scanning for viruses

The **malware** ACL condition lets you connect virus scanner software to Exim. It supports a “generic” interface to scanners called via the shell, and specialized interfaces for “daemon” type virus scanners, which are resident in memory and thus are much faster.

A timeout of 2 minutes is applied to a scanner call (by default); if it expires then a defer action is taken.

You can set the **av\_scanner** option in the main part of the configuration to specify which scanner to use, together with any additional options that are needed. The basic syntax is as follows:

```
av_scanner = <scanner-type>:<option1>:<option2>:[...]
```

If you do not set **av\_scanner**, it defaults to

```
av_scanner = sophie:/var/run/sophie
```

If the value of **av\_scanner** starts with a dollar character, it is expanded before use. The usual list-parsing of the content (see 6.20) applies. The following scanner types are supported in this release:

#### **avast**

This is the scanner daemon of Avast. It has been tested with Avast Core Security (currently at version 1.1.7). You can get a trial version at <http://www.avast.com> or for Linux at <http://www.avast.com/linux-server-antivirus>. This scanner type takes one option, which can be either a full path to a UNIX socket, or host and port specifiers separated by white space. The host may be a name or an IP address; the port is either a single number or a pair of numbers with a dash between. Any further options are given, on separate lines, to the daemon as options before the main scan command. For example:

```
av_scanner = avast:/var/run/avast/scan.sock:FLAGS -fullfiles:SENSITIVITY -pup
av_scanner = avast:192.168.2.22 5036
```

If you omit the argument, the default path `/var/run/avast/scan.sock` is used. If you use a remote host, you need to make Exim's spool directory available to it, as the scanner is passed a file path, not file contents. For information about available commands and their options you may use

```
$ socat UNIX:/var/run/avast/scan.sock STDIO:
    FLAGS
    SENSITIVITY
    PACK
```

#### **aveserver**

This is the scanner daemon of Kaspersky Version 5. You can get a trial version at <http://www.kaspersky.com>. This scanner type takes one option, which is the path to the daemon's UNIX socket. The default is shown in this example:

```
av_scanner = aveserver:/var/run/aveserver
```

#### **clamd**

This daemon-type scanner is GPL and free. You can get it at <http://www.clamav.net/>. Some older versions of clamd do not seem to unpack MIME containers, so it used to be recommended to unpack MIME attachments in the MIME ACL. This is no longer believed to be necessary.

The options are a list of server specifiers, which may be a UNIX socket specification, a TCP socket specification, or a (global) option.

A socket specification consists of a space-separated list. For a Unix socket the first element is a full path for the socket, for a TCP socket the first element is the IP address and the second a port number. Any further elements are per-server (non-global) options. These per-server options are supported:

```
retry=<timespec>#Retry on connect fail
```

The `retry` option specifies a time after which a single retry for a failed connect is made. The default is to not retry.

If a Unix socket file is specified, only one server is supported.

Examples:

```
av_scanner = clamd:/opt/clamd/socket
av_scanner = clamd:192.0.2.3 1234
av_scanner = clamd:192.0.2.3 1234:local
av_scanner = clamd:192.0.2.3 1234 retry=10s
av_scanner = clamd:192.0.2.3 1234 : 192.0.2.4 1234
```

If the value of `av_scanner` points to a UNIX socket file or contains the `local` option, then the ClamAV interface will pass a filename containing the data to be scanned, which will should normally result in less I/O happening and be more efficient. Normally in the TCP case, the data is streamed to ClamAV as Exim does not assume that there is a common filesystem with the remote host. There is an option `WITH_OLD_CLAMAV_STREAM` in `src/EDITME` available, should you be running a version of ClamAV prior to 0.95.

The final example shows that multiple TCP targets can be specified. Exim will randomly use one for each incoming email (i.e. it load balances them). Note that only TCP targets may be used if specifying a list of scanners; a UNIX socket cannot be mixed in with TCP targets. If one of the servers becomes unavailable, Exim will try the remaining one(s) until it finds one that works. When a clamd server becomes unreachable, Exim will log a message. Exim does not keep track of scanner state between multiple messages, and the scanner selection is random, so the message will get logged in the mainlog for each email that the down scanner gets chosen first (message wrapped to be readable):

```
2013-10-09 14:30:39 1VTumd-0000Y8-BQ malware acl condition:
      clamd: connection to localhost, port 3310 failed
      (Connection refused)
```

If the option is unset, the default is */tmp/clamd*. Thanks to David Saez for contributing the code for this scanner.

### **cmdline**

This is the keyword for the generic command line scanner interface. It can be used to attach virus scanners that are invoked from the shell. This scanner type takes 3 mandatory options:

- (1) The full path and name of the scanner binary, with all command line options, and a placeholder (%s) for the directory to scan.
- (2) A regular expression to match against the STDOUT and STDERR output of the virus scanner. If the expression matches, a virus was found. You must make absolutely sure that this expression matches on “virus found”. This is called the “trigger” expression.
- (3) Another regular expression, containing exactly one pair of parentheses, to match the name of the virus found in the scanners output. This is called the “name” expression.

For example, Sophos Sweep reports a virus on a line like this:

```
Virus 'W32/Magistr-B' found in file ./those.bat
```

For the trigger expression, we can match the phrase “found in file”. For the name expression, we want to extract the W32/Magistr-B string, so we can match for the single quotes left and right of it. Altogether, this makes the configuration setting:

```
av_scanner = cmdline:\
              /path/to/sweep -ss -all -rec -archive %s:\
              found in file:'(.+)'
```

### **drweb**

The DrWeb daemon scanner (<http://www.sald.com/>) interface takes one option, either a full path to a UNIX socket, or host and port specifiers separated by white space. The host may be a name or an IP address; the port is either a single number or a pair of numbers with a dash between. For example:

```
av_scanner = drweb:/var/run/drwebd.sock
av_scanner = drweb:192.168.2.20 31337
```

If you omit the argument, the default path */usr/local/drweb/run/drwebd.sock* is used. Thanks to Alex Miller for contributing the code for this scanner.

### **f-protd**

The f-protd scanner is accessed via HTTP over TCP. One argument is taken, being a space-separated hostname and port number (or port-range). For example:

```
av_scanner = f-protd:localhost 10200-10204
```

If you omit the argument, the default values shown above are used.

### **f-prot6d**

The f-prot6d scanner is accessed using the FPSCAND protocol over TCP. One argument is taken, being a space-separated hostname and port number. For example:

```
av_scanner = f-prot6d:localhost 10200
```

If you omit the argument, the default values show above are used.

### **fsecure**

The F-Secure daemon scanner (<http://www.f-secure.com>) takes one argument which is the path to a UNIX socket. For example:

```
av_scanner = fsecure:/path/to/.fsav
```

If no argument is given, the default is `/var/run/fsav`. Thanks to Johan Thelmen for contributing the code for this scanner.

### **kavdaemon**

This is the scanner daemon of Kaspersky Version 4. This version of the Kaspersky scanner is outdated. Please upgrade (see **aveserver** above). This scanner type takes one option, which is the path to the daemon's UNIX socket. For example:

```
av_scanner = kavdaemon:/opt/AVP/AvpCtl
```

The default path is `/var/run/AvpCtl`.

### **mksd**

This is a daemon type scanner that is aimed mainly at Polish users, though some parts of documentation are now available in English. You can get it at <http://linux.mks.com.pl/>. The only option for this scanner type is the maximum number of processes used simultaneously to scan the attachments, provided that mksd has been run with at least the same number of child processes. For example:

```
av_scanner = mksd:2
```

You can safely omit this option (the default value is 1).

### **sock**

This is a general-purpose way of talking to simple scanner daemons running on the local machine. There are four options: an address (which may be an IP address and port, or the path of a Unix socket), a commandline to send (may include a single `%s` which will be replaced with the path to the mail file to be scanned), an RE to trigger on from the returned data, and an RE to extract `malware_name` from the returned data. For example:

```
av_scanner = sock:127.0.0.1 6001:%s:(SPAM|VIRUS):(.*)$
```

Note that surrounding whitespace is stripped from each option, meaning there is no way to specify a trailing newline. The socket specifier and both regular-expressions are required. Default for the commandline is `%s\n` (note this does have a trailing newline); specify an empty element to get this.

### **sophie**

Sophie is a daemon that uses Sophos' **libsavi** library to scan for viruses. You can get Sophie at <http://www.clanfield.info/sophie/>. The only option for this scanner type is the path to the UNIX socket that Sophie uses for client communication. For example:

```
av_scanner = sophie:/tmp/sophie
```

The default path is `/var/run/sophie`, so if you are using this, you can omit the option.

When **av\_scanner** is correctly set, you can use the **malware** condition in the DATA ACL. **Note:** You cannot use the **malware** condition in the MIME ACL.

The **av\_scanner** option is expanded each time **malware** is called. This makes it possible to use different scanners. See further below for an example. The **malware** condition caches its results, so when you use it multiple times for the same message, the actual scanning process is only carried out once. However, using expandable items in **av\_scanner** disables this caching, in which case each use of the **malware** condition causes a new scan of the message.

The **malware** condition takes a right-hand argument that is expanded before use and taken as a list, slash-separated by default. The first element can then be one of

- “true”, “\*”, or “1”, in which case the message is scanned for viruses. The condition succeeds if a virus was found, and fail otherwise. This is the recommended usage.



- “false” or “0” or an empty string, in which case no scanning is done and the condition fails immediately.
- A regular expression, in which case the message is scanned for viruses. The condition succeeds if a virus is found and its name matches the regular expression. This allows you to take special actions on certain types of virus. Note that “/” characters in the RE must be doubled due to the list-processing, unless the separator is changed (in the usual way).

You can append a `defer_ok` element to the **malware** argument list to accept messages even if there is a problem with the virus scanner. Otherwise, such a problem causes the ACL to defer.

You can append a `tmo=<val>` element to the **malware** argument list to specify a non-default timeout. The default is two minutes. For example:

```
malware = * / defer_ok / tmo=10s
```

A timeout causes the ACL to defer.

When a connection is made to the scanner the expansion variable *\$callout\_address* is set to record the actual address used.

When a virus is found, the condition sets up an expansion variable called *\$malware\_name* that contains the name of the virus. You can use it in a **message** modifier that specifies the error returned to the sender, and/or in logging data.

Beware the interaction of Exim’s **message\_size\_limit** with any size limits imposed by your anti-virus scanner.

Here is a very simple scanning example:

```
deny message = This message contains malware ($malware_name)
malware = *
```

The next example accepts messages when there is a problem with the scanner:

```
deny message = This message contains malware ($malware_name)
malware = */defer_ok
```

The next example shows how to use an ACL variable to scan with both sophie and aveserver. It assumes you have set:

```
av_scanner = $acl_m0
```

in the main Exim configuration.

```
deny message = This message contains malware ($malware_name)
set acl_m0 = sophie
malware = *
```

```
deny message = This message contains malware ($malware_name)
set acl_m0 = aveserver
malware = *
```

## 44.2 Scanning with SpamAssassin and Rspamd

The **spam** ACL condition calls SpamAssassin’s **spamd** daemon to get a spam score and a report for the message. Support is also provided for Rspamd.

For more information about installation and configuration of SpamAssassin or Rspamd refer to their respective websites at <http://spamassassin.apache.org> and <http://www.rspamd.com>

SpamAssassin can be installed with CPAN by running:

```
perl -MCPAN -e 'install Mail::SpamAssassin'
```

SpamAssassin has its own set of configuration files. Please review its documentation to see how you can tweak it. The default installation should work nicely, however.

By default, SpamAssassin listens on 127.0.0.1, TCP port 783 and if you intend to use an instance running on the local host you do not need to set **spamd\_address**. If you intend to use another host or port for SpamAssassin, you must set the **spamd\_address** option in the global part of the Exim configuration as follows (example):

```
spamd_address = 192.168.99.45 387
```

The SpamAssassin protocol relies on a TCP half-close from the client. If your SpamAssassin client side is running a Linux system with an iptables firewall, consider setting **net.netfilter.nf\_conntrack\_tcp\_timeout\_close\_wait** to at least the timeout, Exim uses when waiting for a response from the SpamAssassin server (currently defaulting to 120s). With a lower value the Linux connection tracking may consider your half-closed connection as dead too soon.

To use Rspamd (which by default listens on all local addresses on TCP port 11333) you should add **variant=rsamd** after the address/port pair, for example:

```
spamd_address = 127.0.0.1 11333 variant=rsamd
```

As of version 2.60, **SpamAssassin** also supports communication over UNIX sockets. If you want to use these, supply **spamd\_address** with an absolute file name instead of an address/port pair:

```
spamd_address = /var/run/spamd_socket
```

You can have multiple **spamd** servers to improve scalability. These can reside on other hardware reachable over the network. To specify multiple **spamd** servers, put multiple address/port pairs in the **spamd\_address** option, separated with colons (the separator can be changed in the usual way):

```
spamd_address = 192.168.2.10 783 : \  
                192.168.2.11 783 : \  
                192.168.2.12 783
```

Up to 32 **spamd** servers are supported. When a server fails to respond to the connection attempt, all other servers are tried until one succeeds. If no server responds, the **spam** condition defers.

Unix and TCP socket specifications may be mixed in any order. Each element of the list is a list itself, space-separated by default and changeable in the usual way; take care to not double the separator.

For TCP socket specifications a host name or IP (v4 or v6, but subject to list-separator quoting rules) address can be used, and the port can be one or a dash-separated pair. In the latter case, the range is tried in strict order.

Elements after the first for Unix sockets, or second for TCP socket, are options. The supported options are:

<b>pri=&lt;priority&gt;</b>	Selection priority
<b>weight=&lt;value&gt;</b>	Selection bias
<b>time=&lt;start&gt;-&lt;end&gt;</b>	Use only between these times of day
<b>retry=&lt;timespec&gt;</b>	Retry on connect fail
<b>tmo=&lt;timespec&gt;</b>	Connection time limit
<b>variant=rsamd</b>	Use Rspamd rather than SpamAssassin protocol

The **pri** option specifies a priority for the server within the list, higher values being tried first. The default priority is 1.

The **weight** option specifies a selection bias. Within a priority set servers are queried in a random fashion, weighted by this value. The default value for selection bias is 1.

Time specifications for the **time** option are <hour>.<minute>.<second> in the local time zone; each element being one or more digits. Either the seconds or both minutes and seconds, plus the leading . characters, may be omitted and will be taken as zero.

Timeout specifications for the **retry** and **tmo** options are the usual Exim time interval standard, e.g. 20s or 1m.

The **tmo** option specifies an overall timeout for communication. The default value is two minutes.

The `retry` option specifies a time after which a single retry for a failed connect is made. The default is to not retry.

The `spamd_address` variable is expanded before use if it starts with a dollar sign. In this case, the expansion may return a string that is used as the list so that multiple spamd servers can be the result of an expansion.

When a connection is made to the server the expansion variable `$callout_address` is set to record the actual address used.

## 44.3 Calling SpamAssassin from an Exim ACL

Here is a simple example of the use of the `spam` condition in a DATA ACL:

```
deny message = This message was classified as SPAM
    spam = joe
```

The right-hand side of the `spam` condition specifies a name. This is relevant if you have set up multiple SpamAssassin profiles. If you do not want to scan using a specific profile, but rather use the SpamAssassin system-wide default profile, you can scan for an unknown name, or simply use “nobody”. Rspamd does not use this setting. However, you must put something on the right-hand side.

The name allows you to use per-domain or per-user antispam profiles in principle, but this is not straightforward in practice, because a message may have multiple recipients, not necessarily all in the same domain. Because the `spam` condition has to be called from a DATA-time ACL in order to be able to read the contents of the message, the variables `$local_part` and `$domain` are not set. Careful enforcement of single-recipient messages (e.g. by responding with defer in the recipient ACL for all recipients after the first), or the use of PRDR, are needed to use this feature.

The right-hand side of the `spam` condition is expanded before being used, so you can put lookups or conditions there. When the right-hand side evaluates to “0” or “false”, no scanning is done and the condition fails immediately.

Scanning with SpamAssassin uses a lot of resources. If you scan every message, large ones may cause significant performance degradation. As most spam messages are quite small, it is recommended that you do not scan the big ones. For example:

```
deny message = This message was classified as SPAM
    condition = ${if < {$message_size}{10K}}
    spam = nobody
```

The `spam` condition returns true if the threshold specified in the user’s SpamAssassin profile has been matched or exceeded. If you want to use the `spam` condition for its side effects (see the variables below), you can make it always return “true” by appending `:true` to the username.

When the `spam` condition is run, it sets up a number of expansion variables. Except for `$spam_report`, these variables are saved with the received message so are available for use at delivery time.

### *\$spam\_score*

The spam score of the message, for example “3.4” or “30.5”. This is useful for inclusion in log or reject messages.

### *\$spam\_score\_int*

The spam score of the message, multiplied by ten, as an integer value. For example “34” or “305”. It may appear to disagree with `$spam_score` because `$spam_score` is rounded and `$spam_score_int` is truncated. The integer value is useful for numeric comparisons in conditions.

### *\$spam\_bar*

A string consisting of a number of “+” or “-” characters, representing the integer part of the spam score value. A spam score of 4.4 would have a `$spam_bar` value of “++++”. This is useful for inclusion in warning headers, since MUAs can match on such strings. The maximum length of the spam bar is 50 characters.

### *\$spam\_report*

A multiline text table, containing the full SpamAssassin report for the message. Useful for inclusion in headers or reject messages. This variable is only usable in a DATA-time ACL. Beware that SpamAssassin may return non-ASCII characters, especially when running in country-specific locales, which are not legal unencoded in headers.

### *\$spam\_action*

For SpamAssassin either 'reject' or 'no action' depending on the spam score versus threshold. For Rspamd, the recommended action.

The **spam** condition caches its results unless expansion in `spamd_address` was used. If you call it again with the same user name, it does not scan again, but rather returns the same values as before.

The **spam** condition returns DEFER if there is any error while running the message through SpamAssassin or if the expansion of `spamd_address` failed. If you want to treat DEFER as FAIL (to pass on to the next ACL statement block), append `/defer_ok` to the right-hand side of the spam condition, like this:

```
deny message = This message was classified as SPAM
spam         = joe/defer_ok
```

This causes messages to be accepted even if there is a problem with **spamd**.

Here is a longer, commented example of the use of the **spam** condition:

```
# put headers in all messages (no matter if spam or not)
warn spam = nobody:true
  add_header = X-Spam-Score: $spam_score ($spam_bar)
  add_header = X-Spam-Report: $spam_report

# add second subject line with *SPAM* marker when message
# is over threshold
warn spam = nobody
  add_header = Subject: *SPAM* $h_Subject:

# reject spam at high scores (> 12)
deny message = This message scored $spam_score spam points.
  spam = nobody:true
  condition = ${if >{$spam_score_int}{120}{1}{0}}
```

## 44.4 Scanning MIME parts

The **acl\_smtp\_mime** global option specifies an ACL that is called once for each MIME part of an SMTP message, including multipart types, in the sequence of their position in the message. Similarly, the **acl\_not\_smtp\_mime** option specifies an ACL that is used for the MIME parts of non-SMTP messages. These options may both refer to the same ACL if you want the same processing in both cases.

These ACLs are called (possibly many times) just before the **acl\_smtp\_data** ACL in the case of an SMTP message, or just before the **acl\_not\_smtp** ACL in the case of a non-SMTP message. However, a MIME ACL is called only if the message contains a *Content-Type:* header line. When a call to a MIME ACL does not yield “accept”, ACL processing is aborted and the appropriate result code is sent to the client. In the case of an SMTP message, the **acl\_smtp\_data** ACL is not called when this happens.

You cannot use the **malware** or **spam** conditions in a MIME ACL; these can only be used in the DATA or non-SMTP ACLs. However, you can use the **regex** condition to match against the raw MIME part. You can also use the **mime\_regex** condition to match against the decoded MIME part (see section 44.5).

At the start of a MIME ACL, a number of variables are set from the header information for the relevant MIME part. These are described below. The contents of the MIME part are not by default

decoded into a disk file except for MIME parts whose content-type is “message/rfc822”. If you want to decode a MIME part into a disk file, you can use the **decode** condition. The general syntax is:

```
decode = [ /<path> / ] <filename>
```

The right hand side is expanded before use. After expansion, the value can be:

- (1) “0” or “false”, in which case no decoding is done.
- (2) The string “default”. In that case, the file is put in the temporary “default” directory `<spool_directory>/scan/<message_id>/` with a sequential file name consisting of the message id and a sequence number. The full path and name is available in `$mime_decoded_filename` after decoding.
- (3) A full path name starting with a slash. If the full name is an existing directory, it is used as a replacement for the default directory. The filename is then sequentially assigned. If the path does not exist, it is used as the full path and file name.
- (4) If the string does not start with a slash, it is used as the filename, and the default path is then used.

The **decode** condition normally succeeds. It is only false for syntax errors or unusual circumstances such as memory shortages. You can easily decode a file with its original, proposed filename using

```
decode = $mime_filename
```

However, you should keep in mind that `$mime_filename` might contain anything. If you place files outside of the default path, they are not automatically unlinked.

For RFC822 attachments (these are messages attached to messages, with a content-type of “message/rfc822”), the ACL is called again in the same manner as for the primary message, only that the `$mime_is_rfc822` expansion variable is set (see below). Attached messages are always decoded to disk before being checked, and the files are unlinked once the check is done.

The MIME ACL supports the **regex** and **mime\_regex** conditions. These can be used to match regular expressions against raw and decoded MIME parts, respectively. They are described in section 44.5.

The following list describes all expansion variables that are available in the MIME ACL:

#### *\$mime\_boundary*

If the current part is a multipart (see `$mime_is_multipart`) below, it should have a boundary string, which is stored in this variable. If the current part has no boundary parameter in the *Content-Type*: header, this variable contains the empty string.

#### *\$mime\_charset*

This variable contains the character set identifier, if one was found in the *Content-Type*: header. Examples for charset identifiers are:

```
us-ascii
gb2312 (Chinese)
iso-8859-1
```

Please note that this value is not normalized, so you should do matches case-insensitively.

#### *\$mime\_content\_description*

This variable contains the normalized content of the *Content-Description*: header. It can contain a human-readable description of the parts content. Some implementations repeat the filename for attachments here, but they are usually only used for display purposes.

#### *\$mime\_content\_disposition*

This variable contains the normalized content of the *Content-Disposition*: header. You can expect strings like “attachment” or “inline” here.

#### *\$mime\_content\_id*

This variable contains the normalized content of the *Content-ID*: header. This is a unique ID that can be used to reference a part from another part.

#### *\$mime\_content\_size*

This variable is set only after the **decode** modifier (see above) has been successfully run. It contains the size of the decoded part in kilobytes. The size is always rounded up to full kilobytes, so only a completely empty part has a *\$mime\_content\_size* of zero.

#### *\$mime\_content\_transfer\_encoding*

This variable contains the normalized content of the *Content-transfer-encoding*: header. This is a symbolic name for an encoding type. Typical values are “base64” and “quoted-printable”.

#### *\$mime\_content\_type*

If the MIME part has a *Content-Type*: header, this variable contains its value, lowercased, and without any options (like “name” or “charset”). Here are some examples of popular MIME types, as they may appear in this variable:

```
text/plain
text/html
application/octet-stream
image/jpeg
audio/midi
```

If the MIME part has no *Content-Type*: header, this variable contains the empty string.

#### *\$mime\_decoded\_filename*

This variable is set only after the **decode** modifier (see above) has been successfully run. It contains the full path and file name of the file containing the decoded data.

#### *\$mime\_filename*

This is perhaps the most important of the MIME variables. It contains a proposed filename for an attachment, if one was found in either the *Content-Type*: or *Content-Disposition*: headers. The filename will be RFC2047 or RFC2231 decoded, but no additional sanity checks are done. If no filename was found, this variable contains the empty string.

#### *\$mime\_is\_coverletter*

This variable attempts to differentiate the “cover letter” of an e-mail from attached data. It can be used to clamp down on flashy or unnecessarily encoded content in the cover letter, while not restricting attachments at all.

The variable contains 1 (true) for a MIME part believed to be part of the cover letter, and 0 (false) for an attachment. At present, the algorithm is as follows:

- (1) The outermost MIME part of a message is always a cover letter.
- (2) If a multipart/alternative or multipart/related MIME part is a cover letter, so are all MIME subparts within that multipart.
- (3) If any other multipart is a cover letter, the first subpart is a cover letter, and the rest are attachments.
- (4) All parts contained within an attachment multipart are attachments.

As an example, the following will ban “HTML mail” (including that sent with alternative plain text), while allowing HTML files to be attached. HTML coverletter mail attached to non-HMTL coverletter mail will also be allowed:

```
deny message = HTML mail is not accepted here
!condition = $mime_is_rfc822
condition = $mime_is_coverletter
condition = ${if eq{$mime_content_type}{text/html}{1}{0}}
```

#### *\$mime\_is\_multipart*

This variable has the value 1 (true) when the current part has the main type “multipart”, for example “multipart/alternative” or “multipart/mixed”. Since multipart entities only serve as containers for other parts, you may not want to carry out specific actions on them.

*\$mime\_is\_rfc822*

This variable has the value 1 (true) if the current part is not a part of the checked message itself, but part of an attached message. Attached message decoding is fully recursive.

*\$mime\_part\_count*

This variable is a counter that is raised for each processed MIME part. It starts at zero for the very first part (which is usually a multipart). The counter is per-message, so it is reset when processing RFC822 attachments (see *\$mime\_is\_rfc822*). The counter stays set after **acl\_smtp\_mime** is complete, so you can use it in the DATA ACL to determine the number of MIME parts of a message. For non-MIME messages, this variable contains the value -1.

## 44.5 Scanning with regular expressions

You can specify your own custom regular expression matches on the full body of the message, or on individual MIME parts.

The **regex** condition takes one or more regular expressions as arguments and matches them against the full message (when called in the DATA ACL) or a raw MIME part (when called in the MIME ACL). The **regex** condition matches linewise, with a maximum line length of 32K characters. That means you cannot have multiline matches with the **regex** condition.

The **mime\_regex** condition can be called only in the MIME ACL. It matches up to 32K of decoded content (the whole content at once, not linewise). If the part has not been decoded with the **decode** modifier earlier in the ACL, it is decoded automatically when **mime\_regex** is executed (using default path and filename values). If the decoded data is larger than 32K, only the first 32K characters are checked.

The regular expressions are passed as a colon-separated list. To include a literal colon, you must double it. Since the whole right-hand side string is expanded before being used, you must also escape dollar signs and backslashes with more backslashes, or use the \N facility to disable expansion. Here is a simple example that contains two regular expressions:

```
deny message = contains blacklisted regex ($regex_match_string)
  regex = [Mm]ortgage : URGENT BUSINESS PROPOSAL
```

The condition returns true if any one of the regular expressions matches. The *\$regex\_match\_string* expansion variable is then set up and contains the matching regular expression. The expansion variables *\$regex1* *\$regex2* etc are set to any substrings captured by the regular expression.

**Warning:** With large messages, these conditions can be fairly CPU-intensive.

## 45. Adding a local scan function to Exim

In these days of email worms, viruses, and ever-increasing spam, some sites want to apply a lot of checking to messages before accepting them.

The content scanning extension (chapter 44) has facilities for passing messages to external virus and spam scanning software. You can also do a certain amount in Exim itself through string expansions and the **condition** condition in the ACL that runs after the SMTP DATA command or the ACL for non-SMTP messages (see chapter 43), but this has its limitations.

To allow for further customization to a site's own requirements, there is the possibility of linking Exim with a private message scanning function, written in C. If you want to run code that is written in something other than C, you can of course use a little C stub to call it.

The local scan function is run once for every incoming message, at the point when Exim is just about to accept the message. It can therefore be used to control non-SMTP messages from local processes as well as messages arriving via SMTP.

Exim applies a timeout to calls of the local scan function, and there is an option called **local\_scan\_timeout** for setting it. The default is 5 minutes. Zero means "no timeout". Exim also sets up signal handlers for SIGSEGV, SIGILL, SIGFPE, and SIGBUS before calling the local scan function, so that the most common types of crash are caught. If the timeout is exceeded or one of those signals is caught, the incoming message is rejected with a temporary error if it is an SMTP message. For a non-SMTP message, the message is dropped and Exim ends with a non-zero code. The incident is logged on the main and reject logs.

### 45.1 Building Exim to use a local scan function

To make use of the local scan function feature, you must tell Exim where your function is before building Exim, by setting LOCAL\_SCAN\_SOURCE in your *Local/Makefile*. A recommended place to put it is in the *Local* directory, so you might set

```
LOCAL_SCAN_SOURCE=Local/local_scan.c
```

for example. The function must be called *local\_scan()*. It is called by Exim after it has received a message, when the success return code is about to be sent. This is after all the ACLs have been run. The return code from your function controls whether the message is actually accepted or not. There is a commented template function (that just accepts the message) in the file *\_src/local\_scan.c\_*.

If you want to make use of Exim's run time configuration file to set options for your *local\_scan()* function, you must also set

```
LOCAL_SCAN_HAS_OPTIONS=yes
```

in *Local/Makefile* (see section 45.3 below).

### 45.2 API for local\_scan()

You must include this line near the start of your code:

```
#include "local_scan.h"
```

This header file defines a number of variables and other values, and the prototype for the function itself. Exim is coded to use unsigned char values almost exclusively, and one of the things this header defines is a shorthand for unsigned char called *uschar*. It also contains the following macro definitions, to simplify casting character strings and pointers to character strings:

```
#define CS      (char *)
#define CCS     (const char *)
#define CSS     (char **)
#define US      (unsigned char *)
#define CUS     (const unsigned char *)
#define USS     (unsigned char **)
```



The function prototype for *local\_scan()* is:

```
extern int local_scan(int fd, uschar **return_text);
```

The arguments are as follows:

- **fd** is a file descriptor for the file that contains the body of the message (the -D file). The file is open for reading and writing, but updating it is not recommended. **Warning:** You must *not* close this file descriptor.

The descriptor is positioned at character 19 of the file, which is the first character of the body itself, because the first 19 characters are the message id followed by -D and a newline. If you rewind the file, you should use the macro `SPOOL_DATA_START_OFFSET` to reset to the start of the data, just in case this changes in some future version.

- **return\_text** is an address which you can use to return a pointer to a text string at the end of the function. The value it points to on entry is NULL.

The function must return an **int** value which is one of the following macros:

`LOCAL_SCAN_ACCEPT`

The message is accepted. If you pass back a string of text, it is saved with the message, and made available in the variable *\$local\_scan\_data*. No newlines are permitted (if there are any, they are turned into spaces) and the maximum length of text is 1000 characters.

`LOCAL_SCAN_ACCEPT_FREEZE`

This behaves as `LOCAL_SCAN_ACCEPT`, except that the accepted message is queued without immediate delivery, and is frozen.

`LOCAL_SCAN_ACCEPT_QUEUE`

This behaves as `LOCAL_SCAN_ACCEPT`, except that the accepted message is queued without immediate delivery.

`LOCAL_SCAN_REJECT`

The message is rejected; the returned text is used as an error message which is passed back to the sender and which is also logged. Newlines are permitted – they cause a multiline response for SMTP rejections, but are converted to `\n` in log lines. If no message is given, “Administrative prohibition” is used.

`LOCAL_SCAN_TEMPREJECT`

The message is temporarily rejected; the returned text is used as an error message as for `LOCAL_SCAN_REJECT`. If no message is given, “Temporary local problem” is used.

`LOCAL_SCAN_REJECT_NOLOGHDR`

This behaves as `LOCAL_SCAN_REJECT`, except that the header of the rejected message is not written to the reject log. It has the effect of unsetting the **rejected\_header** log selector for just this rejection. If **rejected\_header** is already unset (see the discussion of the **log\_selection** option in section 52.15), this code is the same as `LOCAL_SCAN_REJECT`.

`LOCAL_SCAN_TEMPREJECT_NOLOGHDR`

This code is a variation of `LOCAL_SCAN_TEMPREJECT` in the same way that `LOCAL_SCAN_REJECT_NOLOGHDR` is a variation of `LOCAL_SCAN_REJECT`.

If the message is not being received by interactive SMTP, rejections are reported by writing to **stderr** or by sending an email, as configured by the **-oe** command line options.

## 45.3 Configuration options for *local\_scan()*

It is possible to have option settings in the main configuration file that set values in static variables in the *local\_scan()* module. If you want to do this, you must have the line

```
LOCAL_SCAN_HAS_OPTIONS=yes
```

in your *Local/Makefile* when you build Exim. (This line is in *OS/Makefile-Default*, commented out). Then, in the *local\_scan()* source file, you must define static variables to hold the option values, and a table to define them.

The table must be a vector called **local\_scan\_options**, of type `optionlist`. Each entry is a triplet, consisting of a name, an option type, and a pointer to the variable that holds the value. The entries must appear in alphabetical order. Following **local\_scan\_options** you must also define a variable called **local\_scan\_options\_count** that contains the number of entries in the table. Here is a short example, showing two kinds of option:

```
static int my_integer_option = 42;
static uschar *my_string_option = US"a default string";

optionlist local_scan_options[] = {
    { "my_integer", opt_int,      &my_integer_option },
    { "my_string",  opt_stringptr, &my_string_option }
};

int local_scan_options_count =
    sizeof(local_scan_options)/sizeof(optionlist);
```

The values of the variables can now be changed from Exim's runtime configuration file by including a local scan section as in this example:

```
begin local_scan
my_integer = 99
my_string = some string of text...
```

The available types of option data are as follows:

#### **opt\_bool**

This specifies a boolean (true/false) option. The address should point to a variable of type `BOOL`, which will be set to `TRUE` or `FALSE`, which are macros that are defined as "1" and "0", respectively. If you want to detect whether such a variable has been set at all, you can initialize it to `TRUE_UNSET`. (`BOOL` variables are integers underneath, so can hold more than two values.)

#### **opt\_fixed**

This specifies a fixed point number, such as is used for load averages. The address should point to a variable of type `int`. The value is stored multiplied by 1000, so, for example, 1.4142 is truncated and stored as 1414.

#### **opt\_int**

This specifies an integer; the address should point to a variable of type `int`. The value may be specified in any of the integer formats accepted by Exim.

#### **opt\_mkint**

This is the same as **opt\_int**, except that when such a value is output in a **-bP** listing, if it is an exact number of kilobytes or megabytes, it is printed with the suffix K or M.

#### **opt\_octint**

This also specifies an integer, but the value is always interpreted as an octal integer, whether or not it starts with the digit zero, and it is always output in octal.

#### **opt\_stringptr**

This specifies a string value; the address must be a pointer to a variable that points to a string (for example, of type `uschar *`).

#### **opt\_time**

This specifies a time interval value. The address must point to a variable of type `int`. The value that is placed there is a number of seconds.

If the **-bP** command line option is followed by `local_scan`, Exim prints out the values of all the `local_scan()` options.

## **45.4 Available Exim variables**

The header `local_scan.h` gives you access to a number of C variables. These are the only ones that are guaranteed to be maintained from release to release. Note, however, that you can obtain the value of

any Exim expansion variable, including *\$recipients*, by calling *expand\_string()*. The exported C variables are as follows:

### **int body\_linecount**

This variable contains the number of lines in the message's body.

It is not valid if the **spool\_files\_wireformat** option is used.

### **int body\_zerocount**

This variable contains the number of binary zero bytes in the message's body.

It is not valid if the **spool\_files\_wireformat** option is used.

### **unsigned int debug\_selector**

This variable is set to zero when no debugging is taking place. Otherwise, it is a bitmap of debugging selectors. Two bits are identified for use in *local\_scan()*; they are defined as macros:

- The **D\_v** bit is set when **-v** was present on the command line. This is a testing option that is not privileged – any caller may set it. All the other selector bits can be set only by admin users.
- The **D\_local\_scan** bit is provided for use by *local\_scan()*; it is set by the **+local\_scan** debug selector. It is not included in the default set of debugging bits.

Thus, to write to the debugging output only when **+local\_scan** has been selected, you should use code like this:

```
if ((debug_selector & D_local_scan) != 0)
    debug_printf("xxx", ...);
```

### **uschar \*expand\_string\_message**

After a failing call to *expand\_string()* (returned value **NULL**), the variable **expand\_string\_message** contains the error message, zero-terminated.

### **header\_line \*header\_list**

A pointer to a chain of header lines. The **header\_line** structure is discussed below.

### **header\_line \*header\_last**

A pointer to the last of the header lines.

### **uschar \*headers\_charset**

The value of the **headers\_charset** configuration option.

### **BOOL host\_checking**

This variable is **TRUE** during a host checking session that is initiated by the **-bh** command line option.

### **uschar \*interface\_address**

The IP address of the interface that received the message, as a string. This is **NULL** for locally submitted messages.

### **int interface\_port**

The port on which this message was received. When testing with the **-bh** command line option, the value of this variable is **-1** unless a port has been specified via the **-oMi** option.

### **uschar \*message\_id**

This variable contains Exim's message id for the incoming message (the value of *\$message\_exim\_id*) as a zero-terminated string.

### **uschar \*received\_protocol**

The name of the protocol by which the message was received.

### **int recipients\_count**

The number of accepted recipients.

### **recipient\_item \*recipients\_list**

The list of accepted recipients, held in a vector of length **recipients\_count**. The **recipient\_item** structure is discussed below. You can add additional recipients by calling *receive\_add\_recipient()* (see below). You can delete recipients by removing them from the vector and adjusting the value in

**recipients\_count.** In particular, by setting **recipients\_count** to zero you remove all recipients. If you then return the value `LOCAL_SCAN_ACCEPT`, the message is accepted, but immediately blackholed. To replace the recipients, you can set **recipients\_count** to zero and then call `receive_add_recipient()` as often as needed.

**uschar \*sender\_address**

The envelope sender address. For bounce messages this is the empty string.

**uschar \*sender\_host\_address**

The IP address of the sending host, as a string. This is NULL for locally-submitted messages.

**uschar \*sender\_host\_authenticated**

The name of the authentication mechanism that was used, or NULL if the message was not received over an authenticated SMTP connection.

**uschar \*sender\_host\_name**

The name of the sending host, if known.

**int sender\_host\_port**

The port on the sending host.

**BOOL smtp\_input**

This variable is TRUE for all SMTP input, including BSMTP.

**BOOL smtp\_batched\_input**

This variable is TRUE for BSMTP input.

**int store\_pool**

The contents of this variable control which pool of memory is used for new requests. See section 45.8 for details.

## 45.5 Structure of header lines

The **header\_line** structure contains the members listed below. You can add additional header lines by calling the `header_add()` function (see below). You can cause header lines to be ignored (deleted) by setting their type to `*`.

**struct header\_line \*next**

A pointer to the next header line, or NULL for the last line.

**int type**

A code identifying certain headers that Exim recognizes. The codes are printing characters, and are documented in chapter 56 of this manual. Notice in particular that any header line whose type is `*` is not transmitted with the message. This flagging is used for header lines that have been rewritten, or are to be removed (for example, *Envelope-sender:* header lines.) Effectively, `*` means “deleted”.

**int slen**

The number of characters in the header line, including the terminating and any internal newlines.

**uschar \*text**

A pointer to the text of the header. It always ends with a newline, followed by a zero byte. Internal newlines are preserved.

## 45.6 Structure of recipient items

The **recipient\_item** structure contains these members:

**uschar \*address**

This is a pointer to the recipient address as it was received.

**int pno**

This is used in later Exim processing when top level addresses are created by the **one\_time** option. It is not relevant at the time `local_scan()` is run and must always contain -1 at this stage.

## **uschar \*errors\_to**

If this value is not NULL, bounce messages caused by failing to deliver to the recipient are sent to the address it contains. In other words, it overrides the envelope sender for this one recipient. (Compare the **errors\_to** generic router option.) If a *local\_scan()* function sets an **errors\_to** field to an unqualified address, Exim qualifies it using the domain from **qualify\_recipient**. When *local\_scan()* is called, the **errors\_to** field is NULL for all recipients.

## **45.7 Available Exim functions**

The header *local\_scan.h* gives you access to a number of Exim functions. These are the only ones that are guaranteed to be maintained from release to release:

### **pid\_t child\_open(uschar \*\*argv, uschar \*\*envp, int newumask, int \*infdptr, int \*outfdptr, BOOL make\_leader)**

This function creates a child process that runs the command specified by **argv**. The environment for the process is specified by **envp**, which can be NULL if no environment variables are to be passed. A new umask is supplied for the process in **newumask**.

Pipes to the standard input and output of the new process are set up and returned to the caller via the **infdptr** and **outfdptr** arguments. The standard error is cloned to the standard output. If there are any file descriptors “in the way” in the new process, they are closed. If the final argument is TRUE, the new process is made into a process group leader.

The function returns the pid of the new process, or -1 if things go wrong.

### **int child\_close(pid\_t pid, int timeout)**

This function waits for a child process to terminate, or for a timeout (in seconds) to expire. A timeout value of zero means wait as long as it takes. The return value is as follows:

- $\geq 0$

The process terminated by a normal exit and the value is the process ending status.

- $< 0$  and  $> -256$

The process was terminated by a signal and the value is the negation of the signal number.

- -256

The process timed out.

- -257

There was some other error in wait(); **errno** is still set.

### **pid\_t child\_open\_exim(int \*fd)**

This function provides you with a means of submitting a new message to Exim. (Of course, you can also call */usr/sbin/sendmail* yourself if you want, but this packages it all up for you.) The function creates a pipe, forks a subprocess that is running

```
exim -t -oem -oi -f <>
```

and returns to you (via the `int *` argument) a file descriptor for the pipe that is connected to the standard input. The yield of the function is the PID of the subprocess. You can then write a message to the file descriptor, with recipients in *To:*, *Cc:*, and/or *Bcc:* header lines.

When you have finished, call *child\_close()* to wait for the process to finish and to collect its ending status. A timeout value of zero is usually fine in this circumstance. Unless you have made a mistake with the recipient addresses, you should get a return code of zero.

### **pid\_t child\_open\_exim2(int \*fd, uschar \*sender, uschar \*sender\_authentication)**

This function is a more sophisticated version of *child\_open()*. The command that it runs is:

```
exim -t -oem -oi -f sender -oMas sender_authentication
```

The third argument may be NULL, in which case the **-oMas** option is omitted.

**void debug\_printf(char \*, ...)**

This is Exim's debugging function, with arguments as for *printf()*. The output is written to the standard error stream. If no debugging is selected, calls to *debug\_printf()* have no effect. Normally, you should make calls conditional on the *local\_scan* debug selector by coding like this:

```
if ((debug_selector & D_local_scan) != 0)
    debug_printf("xxx", ...);
```

**uschar \*expand\_string(uschar \*string)**

This is an interface to Exim's string expansion code. The return value is the expanded string, or NULL if there was an expansion failure. The C variable **expand\_string\_message** contains an error message after an expansion failure. If expansion does not change the string, the return value is the pointer to the input string. Otherwise, the return value points to a new block of memory that was obtained by a call to *store\_get()*. See section 45.8 below for a discussion of memory handling.

**void header\_add(int type, char \*format, ...)**

This function allows you to add additional header line at the end of the existing ones. The first argument is the type, and should normally be a space character. The second argument is a format string and any number of substitution arguments as for *sprintf()*. You may include internal newlines if you want, and you must ensure that the string ends with a newline.

**void header\_add\_at\_position(BOOL after, uschar \*name, BOOL topnot, int type, char \*format, ...)**

This function adds a new header line at a specified point in the header chain. The header itself is specified as for *header\_add()*.

If **name** is NULL, the new header is added at the end of the chain if **after** is true, or at the start if **after** is false. If **name** is not NULL, the header lines are searched for the first non-deleted header that matches the name. If one is found, the new header is added before it if **after** is false. If **after** is true, the new header is added after the found header and any adjacent subsequent ones with the same name (even if marked "deleted"). If no matching non-deleted header is found, the **topnot** option controls where the header is added. If it is true, addition is at the top; otherwise at the bottom. Thus, to add a header after all the *Received:* headers, or at the top if there are no *Received:* headers, you could use

```
header_add_at_position(TRUE, US"Received", TRUE,
    ' ', "X-xxx: ...");
```

Normally, there is always at least one non-deleted *Received:* header, but there may not be if **received\_header\_text** expands to an empty string.

**void header\_remove(int occurrence, uschar \*name)**

This function removes header lines. If **occurrence** is zero or negative, all occurrences of the header are removed. If occurrence is greater than zero, that particular instance of the header is removed. If no header(s) can be found that match the specification, the function does nothing.

**BOOL header\_testname(header\_line \*hdr, uschar \*name, int length, BOOL notdel)**

This function tests whether the given header has the given name. It is not just a string comparison, because white space is permitted between the name and the colon. If the **notdel** argument is true, a false return is forced for all "deleted" headers; otherwise they are not treated specially. For example:

```
if (header_testname(h, US"X-Spam", 6, TRUE)) ...
```

**uschar \*lss\_b64encode(uschar \*cleartext, int length)**

This function base64-encodes a string, which is passed by address and length. The text may contain bytes of any value, including zero. The result is passed back in dynamic memory that is obtained by calling *store\_get()*. It is zero-terminated.

**int lss\_b64decode(uschar \*codetext, uschar \*\*cleartext)**

This function decodes a base64-encoded string. Its arguments are a zero-terminated base64-encoded string and the address of a variable that is set to point to the result, which is in dynamic memory. The length of the decoded string is the yield of the function. If the input is invalid base64 data, the yield is -1. A zero byte is added to the end of the output string to make it easy to interpret

as a C string (assuming it contains no zeros of its own). The added zero byte is not included in the returned count.

### **int lss\_match\_domain(uschar \*domain, uschar \*list)**

This function checks for a match in a domain list. Domains are always matched caselessly. The return value is one of the following:

OK	match succeeded
FAIL	match failed
DEFER	match deferred

DEFER is usually caused by some kind of lookup defer, such as the inability to contact a database.

### **int lss\_match\_local\_part(uschar \*localpart, uschar \*list, BOOL caseless)**

This function checks for a match in a local part list. The third argument controls case-sensitivity. The return values are as for *lss\_match\_domain()*.

### **int lss\_match\_address(uschar \*address, uschar \*list, BOOL caseless)**

This function checks for a match in an address list. The third argument controls the case-sensitivity of the local part match. The domain is always matched caselessly. The return values are as for *lss\_match\_domain()*.

### **int lss\_match\_host(uschar \*host\_name, uschar \*host\_address, uschar \*list)**

This function checks for a match in a host list. The most common usage is expected to be

```
lss_match_host(sender_host_name, sender_host_address, ...)
```

An empty address field matches an empty item in the host list. If the host name is NULL, the name corresponding to *\$sender\_host\_address* is automatically looked up if a host name is required to match an item in the list. The return values are as for *lss\_match\_domain()*, but in addition, *lss\_match\_host()* returns ERROR in the case when it had to look up a host name, but the lookup failed.

### **void log\_write(unsigned int selector, int which, char \*format, ...)**

This function writes to Exim's log files. The first argument should be zero (it is concerned with **log\_selector**). The second argument can be LOG\_MAIN or LOG\_REJECT or LOG\_PANIC or the inclusive "or" of any combination of them. It specifies to which log or logs the message is written. The remaining arguments are a format and relevant insertion arguments. The string should not contain any newlines, not even at the end.

### **void receive\_add\_recipient(uschar \*address, int pno)**

This function adds an additional recipient to the message. The first argument is the recipient address. If it is unqualified (has no domain), it is qualified with the **qualify\_recipient** domain. The second argument must always be -1.

This function does not allow you to specify a private **errors\_to** address (as described with the structure of **recipient\_item** above), because it pre-dates the addition of that field to the structure. However, it is easy to add such a value afterwards. For example:

```
receive_add_recipient(US"monitor@mydom.example", -1);
recipients_list[recipients_count-1].errors_to =
    US"postmaster@mydom.example";
```

### **BOOL receive\_remove\_recipient(uschar \*recipient)**

This is a convenience function to remove a named recipient from the list of recipients. It returns true if a recipient was removed, and false if no matching recipient could be found. The argument must be a complete email address.

### **uschar rfc2047\_decode(uschar \*string, BOOL lencheck, uschar \*target, int zeroval, int \*lenptr, uschar \*\*error)**

This function decodes strings that are encoded according to RFC 2047. Typically these are the contents of header lines. First, each "encoded word" is decoded from the Q or B encoding into a byte-string. Then, if provided with the name of a charset encoding, and if the *iconv()* function is available, an attempt is made to translate the result to the named character set. If this fails, the binary string is returned with an error message.

The first argument is the string to be decoded. If **lencheck** is TRUE, the maximum MIME word length is enforced. The third argument is the target encoding, or NULL if no translation is wanted.

If a binary zero is encountered in the decoded string, it is replaced by the contents of the **zeroval** argument. For use with Exim headers, the value must not be 0 because header lines are handled as zero-terminated strings.

The function returns the result of processing the string, zero-terminated; if **lenptr** is not NULL, the length of the result is set in the variable to which it points. When **zeroval** is 0, **lenptr** should not be NULL.

If an error is encountered, the function returns NULL and uses the **error** argument to return an error message. The variable pointed to by **error** is set to NULL if there is no error; it may be set non-NULL even when the function returns a non-NULL value if decoding was successful, but there was a problem with translation.

#### **int smtp\_fflush(void)**

This function is used in conjunction with *smtp\_printf()*, as described below.

#### **void smtp\_printf(char \*, ...)**

The arguments of this function are like *printf()*; it writes to the SMTP output stream. You should use this function only when there is an SMTP output stream, that is, when the incoming message is being received via interactive SMTP. This is the case when **smtp\_input** is TRUE and **smtp\_batched\_input** is FALSE. If you want to test for an incoming message from another host (as opposed to a local process that used the **-bs** command line option), you can test the value of **sender\_host\_address**, which is non-NULL when a remote host is involved.

If an SMTP TLS connection is established, *smtp\_printf()* uses the TLS output function, so it can be used for all forms of SMTP connection.

Strings that are written by *smtp\_printf()* from within *local\_scan()* must start with an appropriate response code: 550 if you are going to return LOCAL\_SCAN\_REJECT, 451 if you are going to return LOCAL\_SCAN\_TEMPREJECT, and 250 otherwise. Because you are writing the initial lines of a multi-line response, the code must be followed by a hyphen to indicate that the line is not the final response line. You must also ensure that the lines you write terminate with CRLF. For example:

```
smtp_printf("550-this is some extra info\r\n");
return LOCAL_SCAN_REJECT;
```

Note that you can also create multi-line responses by including newlines in the data returned via the **return\_text** argument. The added value of using *smtp\_printf()* is that, for instance, you could introduce delays between multiple output lines.

The *smtp\_printf()* function does not return any error indication, because it does not automatically flush pending output, and therefore does not test the state of the stream. (In the main code of Exim, flushing and error detection is done when Exim is ready for the next SMTP input command.) If you want to flush the output and check for an error (for example, the dropping of a TCP/IP connection), you can call *smtp\_fflush()*, which has no arguments. It flushes the output stream, and returns a non-zero value if there is an error.

#### **void \*store\_get(int)**

This function accesses Exim's internal store (memory) manager. It gets a new chunk of memory whose size is given by the argument. Exim bombs out if it ever runs out of memory. See the next section for a discussion of memory handling.

#### **void \*store\_get\_perm(int)**

This function is like *store\_get()*, but it always gets memory from the permanent pool. See the next section for a discussion of memory handling.

#### **uschar \*string\_copy(uschar \*string)**

See below.

#### **uschar \*string\_copyn(uschar \*string, int length)**

See below.



**uschar \*string\_sprintf(char \*format, ...)**

These three functions create strings using Exim's dynamic memory facilities. The first makes a copy of an entire string. The second copies up to a maximum number of characters, indicated by the second argument. The third uses a format and insertion arguments to create a new string. In each case, the result is a pointer to a new string in the current memory pool. See the next section for more discussion.

## 45.8 More about Exim's memory handling

No function is provided for freeing memory, because that is never needed. The dynamic memory that Exim uses when receiving a message is automatically recycled if another message is received by the same process (this applies only to incoming SMTP connections – other input methods can supply only one message at a time). After receiving the last message, a reception process terminates.

Because it is recycled, the normal dynamic memory cannot be used for holding data that must be preserved over a number of incoming messages on the same SMTP connection. However, Exim in fact uses two pools of dynamic memory; the second one is not recycled, and can be used for this purpose.

If you want to allocate memory that remains available for subsequent messages in the same SMTP connection, you should set

```
store_pool = POOL_PERM
```

before calling the function that does the allocation. There is no need to restore the value if you do not need to; however, if you do want to revert to the normal pool, you can either restore the previous value of **store\_pool** or set it explicitly to **POOL\_MAIN**.

The pool setting applies to all functions that get dynamic memory, including *expand\_string()*, *store\_get()*, and the *string\_xxx()* functions. There is also a convenience function called *store\_get\_perm()* that gets a block of memory from the permanent pool while preserving the value of **store\_pool**.

## 46. System-wide message filtering

The previous chapters (on ACLs and the local scan function) describe checks that can be applied to messages before they are accepted by a host. There is also a mechanism for checking messages once they have been received, but before they are delivered. This is called the *system filter*.

The system filter operates in a similar manner to users' filter files, but it is run just once per message (however many recipients the message has). It should not normally be used as a substitute for routing, because **deliver** commands in a system router provide new envelope recipient addresses. The system filter must be an Exim filter. It cannot be a Sieve filter.

The system filter is run at the start of a delivery attempt, before any routing is done. If a message fails to be completely delivered at the first attempt, the system filter is run again at the start of every retry. If you want your filter to do something only once per message, you can make use of the **first\_delivery** condition in an **if** command in the filter to prevent it happening on retries.

**Warning:** Because the system filter runs just once, variables that are specific to individual recipient addresses, such as *\$local\_part* and *\$domain*, are not set, and the “personal” condition is not meaningful. If you want to run a centrally-specified filter for each recipient address independently, you can do so by setting up a suitable *redirect* router, as described in section 46.8 below.

### 46.1 Specifying a system filter

The name of the file that contains the system filter must be specified by setting **system\_filter**. If you want the filter to run under a uid and gid other than root, you must also set **system\_filter\_user** and **system\_filter\_group** as appropriate. For example:

```
system_filter = /etc/mail/exim.filter
system_filter_user = exim
```

If a system filter generates any deliveries directly to files or pipes (via the **save** or **pipe** commands), transports to handle these deliveries must be specified by setting **system\_filter\_file\_transport** and **system\_filter\_pipe\_transport**, respectively. Similarly, **system\_filter\_reply\_transport** must be set to handle any messages generated by the **reply** command.

### 46.2 Testing a system filter

You can run simple tests of a system filter in the same way as for a user filter, but you should use **-bF** rather than **-bf**, so that features that are permitted only in system filters are recognized.

If you want to test the combined effect of a system filter and a user filter, you can use both **-bF** and **-bf** on the same command line.

### 46.3 Contents of a system filter

The language used to specify system filters is the same as for users' filter files. It is described in the separate end-user document *Exim's interface to mail filtering*. However, there are some additional features that are available only in system filters; these are described in subsequent sections. If they are encountered in a user's filter file or when testing with **-bf**, they cause errors.

There are two special conditions which, though available in users' filter files, are designed for use in system filters. The condition **first\_delivery** is true only for the first attempt at delivering a message, and **manually\_thawed** is true only if the message has been frozen, and subsequently thawed by an admin user. An explicit forced delivery counts as a manual thaw, but thawing as a result of the **auto\_thaw** setting does not.

**Warning:** If a system filter uses the **first\_delivery** condition to specify an “unseen” (non-significant) delivery, and that delivery does not succeed, it will not be tried again. If you want Exim to retry an unseen delivery until it succeeds, you should arrange to set it up every time the filter runs.

When a system filter finishes running, the values of the variables *\$n0* – *\$n9* are copied into *\$sn0* – *\$sn9* and are thereby made available to users' filter files. Thus a system filter can, for example, set up "scores" to which users' filter files can refer.

## 46.4 Additional variable for system filters

The expansion variable *\$recipients*, containing a list of all the recipients of the message (separated by commas and white space), is available in system filters. It is not available in users' filters for privacy reasons.

## 46.5 Defer, freeze, and fail commands for system filters

There are three extra commands (**defer**, **freeze** and **fail**) which are always available in system filters, but are not normally enabled in users' filters. (See the **allow\_defer**, **allow\_freeze** and **allow\_fail** options for the *redirect* router.) These commands can optionally be followed by the word **text** and a string containing an error message, for example:

```
fail text "this message looks like spam to me"
```

The keyword **text** is optional if the next character is a double quote.

The **defer** command defers delivery of the original recipients of the message. The **fail** command causes all the original recipients to be failed, and a bounce message to be created. The **freeze** command suspends all delivery attempts for the original recipients. In all cases, any new deliveries that are specified by the filter are attempted as normal after the filter has run.

The **freeze** command is ignored if the message has been manually unfrozen and not manually frozen since. This means that automatic freezing by a system filter can be used as a way of checking out suspicious messages. If a message is found to be all right, manually unfreezing it allows it to be delivered.

The text given with a fail command is used as part of the bounce message as well as being written to the log. If the message is quite long, this can fill up a lot of log space when such failures are common. To reduce the size of the log message, Exim interprets the text in a special way if it starts with the two characters << and contains >> later. The text between these two strings is written to the log, and the rest of the text is used in the bounce message. For example:

```
fail "<<filter test 1>>Your message is rejected \
    because it contains attachments that we are \
    not prepared to receive."
```

Take great care with the **fail** command when basing the decision to fail on the contents of the message, because the bounce message will of course include the contents of the original message and will therefore trigger the **fail** command again (causing a mail loop) unless steps are taken to prevent this. Testing the **error\_message** condition is one way to prevent this. You could use, for example

```
if $message_body contains "this is spam" and not error_message
then fail text "spam is not wanted here" endif
```

though of course that might let through unwanted bounce messages. The alternative is clever checking of the body and/or headers to detect bounces generated by the filter.

The interpretation of a system filter file ceases after a **defer**, **freeze**, or **fail** command is obeyed. However, any deliveries that were set up earlier in the filter file are honoured, so you can use a sequence such as

```
mail ...
freeze
```

to send a specified message when the system filter is freezing (or deferring or failing) a message. The normal deliveries for the message do not, of course, take place.

## 46.6 Adding and removing headers in a system filter

Two filter commands that are available only in system filters are:

```
headers add <string>
headers remove <string>
```

The argument for the **headers add** is a string that is expanded and then added to the end of the message's headers. It is the responsibility of the filter maintainer to make sure it conforms to RFC 2822 syntax. Leading white space is ignored, and if the string is otherwise empty, or if the expansion is forced to fail, the command has no effect.

You can use “\n” within the string, followed by white space, to specify continued header lines. More than one header may be added in one command by including “\n” within the string without any following white space. For example:

```
headers add "X-header-1: ....\n \
            continuation of X-header-1 ...\n\
X-header-2: ...."
```

Note that the header line continuation white space after the first newline must be placed before the backslash that continues the input string, because white space after input continuations is ignored.

The argument for **headers remove** is a colon-separated list of header names. This command applies only to those headers that are stored with the message; those that are added at delivery time (such as *Envelope-To:* and *Return-Path:*) cannot be removed by this means. If there is more than one header with the same name, they are all removed.

The **headers** command in a system filter makes an immediate change to the set of header lines that was received with the message (with possible additions from ACL processing). Subsequent commands in the system filter operate on the modified set, which also forms the basis for subsequent message delivery. Unless further modified during routing or transporting, this set of headers is used for all recipients of the message.

During routing and transporting, the variables that refer to the contents of header lines refer only to those lines that are in this set. Thus, header lines that are added by a system filter are visible to users' filter files and to all routers and transports. This contrasts with the manipulation of header lines by routers and transports, which is not immediate, but which instead is saved up until the message is actually being written (see section 47.17).

If the message is not delivered at the first attempt, header lines that were added by the system filter are stored with the message, and so are still present at the next delivery attempt. Header lines that were removed are still present, but marked “deleted” so that they are not transported with the message. For this reason, it is usual to make the **headers** command conditional on **first\_delivery** so that the set of header lines is not modified more than once.

Because header modification in a system filter acts immediately, you have to use an indirect approach if you want to modify the contents of a header line. For example:

```
headers add "Old-Subject: $h_subject:"
headers remove "Subject"
headers add "Subject: new subject (was: $h_old-subject:)"
headers remove "Old-Subject"
```

## 46.7 Setting an errors address in a system filter

In a system filter, if a **deliver** command is followed by

```
errors_to <some address>
```

in order to change the envelope sender (and hence the error reporting) for that delivery, any address may be specified. (In a user filter, only the current user's address can be set.) For example, if some mail is being monitored, you might use

```
unseen deliver monitor@spying.example errors_to root@local.example
```

to take a copy which would not be sent back to the normal error reporting address if its delivery failed.

## 46.8 Per-address filtering

In contrast to the system filter, which is run just once per message for each delivery attempt, it is also possible to set up a system-wide filtering operation that runs once for each recipient address. In this case, variables such as *\$local\_part* and *\$domain* can be used, and indeed, the choice of filter file could be made dependent on them. This is an example of a router which implements such a filter:

```
central_filter:
  check_local_user
  driver = redirect
  domains = +local_domains
  file = /central/filters/$local_part
  no_verify
  allow_filter
  allow_freeze
```

The filter is run in a separate process under its own uid. Therefore, either **check\_local\_user** must be set (as above), in which case the filter is run as the local user, or the **user** option must be used to specify which user to use. If both are set, **user** overrides.

Care should be taken to ensure that none of the commands in the filter file specify a significant delivery if the message is to go on to be delivered to its intended recipient. The router will not then claim to have dealt with the address, so it will be passed on to subsequent routers to be delivered in the normal way.

## 47. Message processing

Exim performs various transformations on the sender and recipient addresses of all messages that it handles, and also on the messages' header lines. Some of these are optional and configurable, while others always take place. All of this processing, except rewriting as a result of routing, and the addition or removal of header lines while delivering, happens when a message is received, before it is placed on Exim's queue.

Some of the automatic processing takes place by default only for “locally-originated” messages. This adjective is used to describe messages that are not received over TCP/IP, but instead are passed to an Exim process on its standard input. This includes the interactive “local SMTP” case that is set up by the **-bs** command line option.

**Note:** Messages received over TCP/IP on the loopback interface (127.0.0.1 or ::1) are not considered to be locally-originated. Exim does not treat the loopback interface specially in any way.

If you want the loopback interface to be treated specially, you must ensure that there are appropriate entries in your ACLs.

### 47.1 Submission mode for non-local messages

Processing that happens automatically for locally-originated messages (unless **suppress\_local\_fixups** is set) can also be requested for messages that are received over TCP/IP. The term “submission mode” is used to describe this state. Submission mode is set by the modifier

```
control = submission
```

in a MAIL, RCPT, or pre-data ACL for an incoming message (see sections 43.21 and 43.22). This makes Exim treat the message as a local submission, and is normally used when the source of the message is known to be an MUA running on a client host (as opposed to an MTA). For example, to set submission mode for messages originating on the IPv4 loopback interface, you could include the following in the MAIL ACL:

```
warn hosts = 127.0.0.1
    control = submission
```

There are some options that can be used when setting submission mode. A slash is used to separate options. For example:

```
control = submission/sender_retain
```

Specifying **sender\_retain** has the effect of setting **local\_sender\_retain** true and **local\_from\_check** false for the current incoming message. The first of these allows an existing *Sender:* header in the message to remain, and the second suppresses the check to ensure that *From:* matches the authenticated sender. With this setting, Exim still fixes up messages by adding *Date:* and *Message-ID:* header lines if they are missing, but makes no attempt to check sender authenticity in header lines.

When **sender\_retain** is not set, a submission mode setting may specify a domain to be used when generating a *From:* or *Sender:* header line. For example:

```
control = submission/domain=some.domain
```

The domain may be empty. How this value is used is described in sections 47.11 and 47.16. There is also a **name** option that allows you to specify the user's full name for inclusion in a created *Sender:* or *From:* header line. For example:

```
accept authenticated = *
    control = submission/domain=wonderland.example/\
              name=${lookup {$authenticated_id} \
              lsearch {/etc/exim/namelist}}
```

Because the name may contain any characters, including slashes, the **name** option must be given last. The remainder of the string is used as the name. For the example above, if `/etc/exim/namelist` contains:

```
bigegg: Humpty Dumpty
```

then when the sender has authenticated as *bigegg*, the generated *Sender:* line would be:

```
Sender: Humpty Dumpty <bigegg@wonderland.example>
```

By default, submission mode forces the return path to the same address as is used to create the *Sender:* header. However, if **sender\_retain** is specified, the return path is also left unchanged.

**Note:** The changes caused by submission mode take effect after the predata ACL. This means that any sender checks performed before the fix-ups use the untrusted sender address specified by the user, not the trusted sender address specified by submission mode. Although this might be slightly unexpected, it does mean that you can configure ACL checks to spot that a user is trying to spoof another's address.

## 47.2 Line endings

RFC 2821 specifies that CRLF (two characters: carriage-return, followed by linefeed) is the line ending for messages transmitted over the Internet using SMTP over TCP/IP. However, within individual operating systems, different conventions are used. For example, Unix-like systems use just LF, but others use CRLF or just CR.

Exim was designed for Unix-like systems, and internally, it stores messages using the system's convention of a single LF as a line terminator. When receiving a message, all line endings are translated to this standard format. Originally, it was thought that programs that passed messages directly to an MTA within an operating system would use that system's convention. Experience has shown that this is not the case; for example, there are Unix applications that use CRLF in this circumstance. For this reason, and for compatibility with other MTAs, the way Exim handles line endings for all messages is now as follows:

- LF not preceded by CR is treated as a line ending.
- CR is treated as a line ending; if it is immediately followed by LF, the LF is ignored.
- The sequence "CR, dot, CR" does not terminate an incoming SMTP message, nor a local message in the state where a line containing only a dot is a terminator.
- If a bare CR is encountered within a header line, an extra space is added after the line terminator so as not to end the header line. The reasoning behind this is that bare CRs in header lines are most likely either to be mistakes, or people trying to play silly games.
- If the first header line received in a message ends with CRLF, a subsequent bare LF in a header line is treated in the same way as a bare CR in a header line.

## 47.3 Unqualified addresses

By default, Exim expects every envelope address it receives from an external host to be fully qualified. Unqualified addresses cause negative responses to SMTP commands. However, because SMTP is used as a means of transporting messages from MUAs running on personal workstations, there is sometimes a requirement to accept unqualified addresses from specific hosts or IP networks.

Exim has two options that separately control which hosts may send unqualified sender or recipient addresses in SMTP commands, namely **sender\_unqualified\_hosts** and **recipient\_unqualified\_hosts**. In both cases, if an unqualified address is accepted, it is qualified by adding the value of **qualify\_domain** or **qualify\_recipient**, as appropriate.

Unqualified addresses in header lines are automatically qualified for messages that are locally originated, unless the **-bnq** option is given on the command line. For messages received over SMTP, unqualified addresses in header lines are qualified only if unqualified addresses are permitted in

SMTP commands. In other words, such qualification is also controlled by **sender\_unqualified\_hosts** and **recipient\_unqualified\_hosts**,

## 47.4 The UUCP From line

Messages that have come from UUCP (and some other applications) often begin with a line containing the envelope sender and a timestamp, following the word “From”. Examples of two common formats are:

```
From a.oakley@berlin.mus Fri Jan 5 12:35 GMT 1996
From f.butler@berlin.mus Fri, 7 Jan 97 14:00:00 GMT
```

This line precedes the RFC 2822 header lines. For compatibility with Sendmail, Exim recognizes such lines at the start of messages that are submitted to it via the command line (that is, on the standard input). It does not recognize such lines in incoming SMTP messages, unless the sending host matches **ignore\_fromline\_hosts** or the **-bs** option was used for a local message and **ignore\_fromline\_local** is set. The recognition is controlled by a regular expression that is defined by the **uucp\_from\_pattern** option, whose default value matches the two common cases shown above and puts the address that follows “From” into *\$1*.

When the caller of Exim for a non-SMTP message that contains a “From” line is a trusted user, the message’s sender address is constructed by expanding the contents of **uucp\_sender\_address**, whose default value is “*\$1*”. This is then parsed as an RFC 2822 address. If there is no domain, the local part is qualified with **qualify\_domain** unless it is the empty string. However, if the command line **-f** option is used, it overrides the “From” line.

If the caller of Exim is not trusted, the “From” line is recognized, but the sender address is not changed. This is also the case for incoming SMTP messages that are permitted to contain “From” lines.

Only one “From” line is recognized. If there is more than one, the second is treated as a data line that starts the body of the message, as it is not valid as a header line. This also happens if a “From” line is present in an incoming SMTP message from a source that is not permitted to send them.

## 47.5 Resent- header lines

RFC 2822 makes provision for sets of header lines starting with the string **Resent-** to be added to a message when it is resent by the original recipient to somebody else. These headers are *Resent-Date:*, *Resent-From:*, *Resent-Sender:*, *Resent-To:*, *Resent-Cc:*, *Resent-Bcc:* and *Resent-Message-Id:*. The RFC says:

*Resent fields are strictly informational. They MUST NOT be used in the normal processing of replies or other such automatic actions on messages.*

This leaves things a bit vague as far as other processing actions such as address rewriting are concerned. Exim treats **Resent-** header lines as follows:

- A *Resent-From:* line that just contains the login id of the submitting user is automatically rewritten in the same way as *From:* (see below).
- If there’s a rewriting rule for a particular header line, it is also applied to **Resent-** header lines of the same type. For example, a rule that rewrites *From:* also rewrites *Resent-From:*.
- For local messages, if *Sender:* is removed on input, *Resent-Sender:* is also removed.
- For a locally-submitted message, if there are any **Resent-** header lines but no *Resent-Date:*, *Resent-From:*, or *Resent-Message-Id:*, they are added as necessary. It is the contents of *Resent-Message-Id:* (rather than *Message-Id:*) which are included in log lines in this case.
- The logic for adding *Sender:* is duplicated for *Resent-Sender:* when any **Resent-** header lines are present.



## 47.6 The Auto-Submitted: header line

Whenever Exim generates an autoreply, a bounce, or a delay warning message, it includes the header line:

```
Auto-Submitted: auto-replied
```

## 47.7 The Bcc: header line

If Exim is called with the **-t** option, to take recipient addresses from a message's header, it removes any *Bcc:* header line that may exist (after extracting its addresses). If **-t** is not present on the command line, any existing *Bcc:* is not removed.

## 47.8 The Date: header line

If a locally-generated or submission-mode message has no *Date:* header line, Exim adds one, using the current date and time, unless the **suppress\_local\_fixups** control has been specified.

## 47.9 The Delivery-date: header line

*Delivery-date:* header lines are not part of the standard RFC 2822 header set. Exim can be configured to add them to the final delivery of messages. (See the generic **delivery\_date\_add** transport option.) They should not be present in messages in transit. If the **delivery\_date\_remove** configuration option is set (the default), Exim removes *Delivery-date:* header lines from incoming messages.

## 47.10 The Envelope-to: header line

*Envelope-to:* header lines are not part of the standard RFC 2822 header set. Exim can be configured to add them to the final delivery of messages. (See the generic **envelope\_to\_add** transport option.) They should not be present in messages in transit. If the **envelope\_to\_remove** configuration option is set (the default), Exim removes *Envelope-to:* header lines from incoming messages.

## 47.11 The From: header line

If a submission-mode message does not contain a *From:* header line, Exim adds one if either of the following conditions is true:

- The envelope sender address is not empty (that is, this is not a bounce message). The added header line copies the envelope sender address.
- The SMTP session is authenticated and *\$authenticated\_id* is not empty.
  - (1) If no domain is specified by the submission control, the local part is *\$authenticated\_id* and the domain is *\$qualify\_domain*.
  - (2) If a non-empty domain is specified by the submission control, the local part is *\$authenticated\_id*, and the domain is the specified domain.
  - (3) If an empty domain is specified by the submission control, *\$authenticated\_id* is assumed to be the complete address.

A non-empty envelope sender takes precedence.

If a locally-generated incoming message does not contain a *From:* header line, and the **suppress\_local\_fixups** control is not set, Exim adds one containing the sender's address. The calling user's login name and full name are used to construct the address, as described in section 47.18. They are obtained from the password data by calling *getpwuid()* (but see the **unknown\_login** configuration option). The address is qualified with **qualify\_domain**.

For compatibility with Sendmail, if an incoming, non-SMTP message has a *From:* header line containing just the unqualified login name of the calling user, this is replaced by an address containing the user's login name and full name as described in section 47.18.

## 47.12 The Message-ID: header line

If a locally-generated or submission-mode incoming message does not contain a *Message-ID:* or *Resent-Message-ID:* header line, and the **suppress\_local\_fixups** control is not set, Exim adds a suitable header line to the message. If there are any *Resent-:* headers in the message, it creates *Resent-Message-ID:*. The id is constructed from Exim's internal message id, preceded by the letter E to ensure it starts with a letter, and followed by @ and the primary host name. Additional information can be included in this header line by setting the **message\_id\_header\_text** and/or **message\_id\_header\_domain** options.

## 47.13 The Received: header line

A *Received:* header line is added at the start of every message. The contents are defined by the **received\_header\_text** configuration option, and Exim automatically adds a semicolon and a timestamp to the configured string.

The *Received:* header is generated as soon as the message's header lines have been received. At this stage, the timestamp in the *Received:* header line is the time that the message started to be received. This is the value that is seen by the DATA ACL and by the *local\_scan()* function.

Once a message is accepted, the timestamp in the *Received:* header line is changed to the time of acceptance, which is (apart from a small delay while the -H spool file is written) the earliest time at which delivery could start.

## 47.14 The References: header line

Messages created by the *autoreply* transport include a *References:* header line. This is constructed according to the rules that are described in section 3.64 of RFC 2822 (which states that replies should contain such a header line), and section 3.14 of RFC 3834 (which states that automatic responses are not different in this respect). However, because some mail processing software does not cope well with very long header lines, no more than 12 message IDs are copied from the *References:* header line in the incoming message. If there are more than 12, the first one and then the final 11 are copied, before adding the message ID of the incoming message.

## 47.15 The Return-path: header line

*Return-path:* header lines are defined as something an MTA may insert when it does the final delivery of messages. (See the generic **return\_path\_add** transport option.) Therefore, they should not be present in messages in transit. If the **return\_path\_remove** configuration option is set (the default), Exim removes *Return-path:* header lines from incoming messages.

## 47.16 The Sender: header line

For a locally-originated message from an untrusted user, Exim may remove an existing *Sender:* header line, and it may add a new one. You can modify these actions by setting the **local\_sender\_retain** option true, the **local\_from\_check** option false, or by using the **suppress\_local\_fixups** control setting.

When a local message is received from an untrusted user and **local\_from\_check** is true (the default), and the **suppress\_local\_fixups** control has not been set, a check is made to see if the address given in the *From:* header line is the correct (local) sender of the message. The address that is expected has the login name as the local part and the value of **qualify\_domain** as the domain. Prefixes and suffixes for the local part can be permitted by setting **local\_from\_prefix** and **local\_from\_suffix** appropriately. If *From:* does not contain the correct sender, a *Sender:* line is added to the message.

If you set **local\_from\_check** false, this checking does not occur. However, the removal of an existing *Sender:* line still happens, unless you also set **local\_sender\_retain** to be true. It is not possible to set both of these options true at the same time.

By default, no processing of *Sender:* header lines is done for messages received over TCP/IP or for messages submitted by trusted users. However, when a message is received over TCP/IP in sub-

mission mode, and **sender\_retain** is not specified on the submission control, the following processing takes place:

First, any existing *Sender:* lines are removed. Then, if the SMTP session is authenticated, and *\$authenticated\_id* is not empty, a sender address is created as follows:

- If no domain is specified by the submission control, the local part is *\$authenticated\_id* and the domain is *\$qualify\_domain*.
- If a non-empty domain is specified by the submission control, the local part is *\$authenticated\_id*, and the domain is the specified domain.
- If an empty domain is specified by the submission control, *\$authenticated\_id* is assumed to be the complete address.

This address is compared with the address in the *From:* header line. If they are different, a *Sender:* header line containing the created address is added. Prefixes and suffixes for the local part in *From:* can be permitted by setting **local\_from\_prefix** and **local\_from\_suffix** appropriately.

**Note:** Whenever a *Sender:* header line is created, the return path for the message (the envelope sender address) is changed to be the same address, except in the case of submission mode when **sender\_retain** is specified.

## 47.17 Adding and removing header lines in routers and transports

When a message is delivered, the addition and removal of header lines can be specified in a system filter, or on any of the routers and transports that process the message. Section 46.6 contains details about modifying headers in a system filter. Header lines can also be added in an ACL as a message is received (see section 43.24).

In contrast to what happens in a system filter, header modifications that are specified on routers and transports apply only to the particular recipient addresses that are being processed by those routers and transports. These changes do not actually take place until a copy of the message is being transported. Therefore, they do not affect the basic set of header lines, and they do not affect the values of the variables that refer to header lines.

**Note:** In particular, this means that any expansions in the configuration of the transport cannot refer to the modified header lines, because such expansions all occur before the message is actually transported.

For both routers and transports, the argument of a **headers\_add** option must be in the form of one or more RFC 2822 header lines, separated by newlines (coded as “\n”). For example:

```
headers_add = X-added-header: added by $primary_hostname\n\
              X-added-second: another added header line
```

Exim does not check the syntax of these added header lines.

Multiple **headers\_add** options for a single router or transport can be specified; the values will append to a single list of header lines. Each header-line is separately expanded.

The argument of a **headers\_remove** option must consist of a colon-separated list of header names. This is confusing, because header names themselves are often terminated by colons. In this case, the colons are the list separators, not part of the names. For example:

```
headers_remove = return-receipt-to:acknowledge-to
```

Multiple **headers\_remove** options for a single router or transport can be specified; the arguments will append to a single header-names list. Each item is separately expanded. Note that colons in complex expansions which are used to form all or part of a **headers\_remove** list will act as list separators.

When **headers\_add** or **headers\_remove** is specified on a router, items are expanded at routing time, and then associated with all addresses that are accepted by that router, and also with any new addresses that it generates. If an address passes through several routers as a result of aliasing or forwarding, the changes are cumulative.

However, this does not apply to multiple routers that result from the use of the **unseen** option. Any header modifications that were specified by the “unseen” router or its predecessors apply only to the “unseen” delivery.

Addresses that end up with different **headers\_add** or **headers\_remove** settings cannot be delivered together in a batch, so a transport is always dealing with a set of addresses that have the same header-processing requirements.

The transport starts by writing the original set of header lines that arrived with the message, possibly modified by the system filter. As it writes out these lines, it consults the list of header names that were attached to the recipient address(es) by **headers\_remove** options in routers, and it also consults the transport’s own **headers\_remove** option. Header lines whose names are on either of these lists are not written out. If there are multiple instances of any listed header, they are all skipped.

After the remaining original header lines have been written, new header lines that were specified by routers’ **headers\_add** options are written, in the order in which they were attached to the address. These are followed by any header lines specified by the transport’s **headers\_add** option.

This way of handling header line modifications in routers and transports has the following consequences:

- The original set of header lines, possibly modified by the system filter, remains “visible”, in the sense that the *\$header\_xxx* variables refer to it, at all times.
- Header lines that are added by a router’s **headers\_add** option are not accessible by means of the *\$header\_xxx* expansion syntax in subsequent routers or the transport.
- Conversely, header lines that are specified for removal by **headers\_remove** in a router remain visible to subsequent routers and the transport.
- Headers added to an address by **headers\_add** in a router cannot be removed by a later router or by a transport.
- An added header can refer to the contents of an original header that is to be removed, even it has the same name as the added header. For example:

```
headers_remove = subject
headers_add = Subject: new subject (was: $h_subject:)
```

**Warning:** The **headers\_add** and **headers\_remove** options cannot be used for a *redirect* router that has the **one\_time** option set.

## 47.18 Constructed addresses

When Exim constructs a sender address for a locally-generated message, it uses the form

```
<user name> <login@qualify_domain>
```

For example:

```
Zaphod Beeblebrox <zaphod@end.univ.example>
```

The user name is obtained from the **-F** command line option if set, or otherwise by looking up the calling user by *getpwuid()* and extracting the “gecos” field from the password entry. If the “gecos” field contains an ampersand character, this is replaced by the login name with the first letter upper cased, as is conventional in a number of operating systems. See the **gecos\_name** option for a way to tailor the handling of the “gecos” field. The **unknown\_username** option can be used to specify user names in cases when there is no password file entry.

In all cases, the user name is made to conform to RFC 2822 by quoting all or parts of it if necessary. In addition, if it contains any non-printing characters, it is encoded as described in RFC 2047, which defines a way of including non-ASCII characters in header lines. The value of the **headers\_charset** option specifies the name of the encoding that is used (the characters are assumed to be in this encoding). The setting of **print\_topbitchars** controls whether characters with the top bit set (that is, with codes greater than 127) count as printing characters or not.

## 47.19 Case of local parts

RFC 2822 states that the case of letters in the local parts of addresses cannot be assumed to be non-significant. Exim preserves the case of local parts of addresses, but by default it uses a lower-cased form when it is routing, because on most Unix systems, usernames are in lower case and case-insensitive routing is required. However, any particular router can be made to use the original case for local parts by setting the **caseful\_local\_part** generic router option.

If you must have mixed-case user names on your system, the best way to proceed, assuming you want case-independent handling of incoming email, is to set up your first router to convert incoming local parts in your domains to the correct case by means of a file lookup. For example:

```
correct_case:
  driver = redirect
  domains = +local_domains
  data = ${lookup{$local_part}cdb\
          {/etc/usercased.cdb}{$value}fail}\
        @$domain
```

For this router, the local part is forced to lower case by the default action (**caseful\_local\_part** is not set). The lower-cased local part is used to look up a new local part in the correct case. If you then set **caseful\_local\_part** on any subsequent routers which process your domains, they will operate on local parts with the correct case in a case-sensitive manner.

## 47.20 Dots in local parts

RFC 2822 forbids empty components in local parts. That is, an unquoted local part may not begin or end with a dot, nor have two consecutive dots in the middle. However, it seems that many MTAs do not enforce this, so Exim permits empty components for compatibility.

## 47.21 Rewriting addresses

Rewriting of sender and recipient addresses, and addresses in headers, can happen automatically, or as the result of configuration options, as described in chapter 31. The headers that may be affected by this are *Bcc:*, *Cc:*, *From:*, *Reply-To:*, *Sender:*, and *To:*.

Automatic rewriting includes qualification, as mentioned above. The other case in which it can happen is when an incomplete non-local domain is given. The routing process may cause this to be expanded into the full domain name. For example, a header such as

```
To: hare@teaparty
```

might get rewritten as

```
To: hare@teaparty.wonderland.fict.example
```

Rewriting as a result of routing is the one kind of message processing that does not happen at input time, as it cannot be done until the address has been routed.

Strictly, one should not do *any* deliveries of a message until all its addresses have been routed, in case any of the headers get changed as a result of routing. However, doing this in practice would hold up many deliveries for unreasonable amounts of time, just because one address could not immediately be routed. Exim therefore does not delay other deliveries when routing of one or more addresses is deferred.

## 48. SMTP processing

Exim supports a number of different ways of using the SMTP protocol, and its LMTP variant, which is an interactive protocol for transferring messages into a closed mail store application. This chapter contains details of how SMTP is processed. For incoming mail, the following are available:

- SMTP over TCP/IP (Exim daemon or *inetd*);
- SMTP over the standard input and output (the **-bs** option);
- Batched SMTP on the standard input (the **-bS** option).

For mail delivery, the following are available:

- SMTP over TCP/IP (the *smtp* transport);
- LMTP over TCP/IP (the *smtp* transport with the **protocol** option set to “lmtp”);
- LMTP over a pipe to a process running in the local host (the *lmtp* transport);
- Batched SMTP to a file or pipe (the *appendfile* and *pipe* transports with the **use\_bsmtp** option set).

*Batched SMTP* is the name for a process in which batches of messages are stored in or read from files (or pipes), in a format in which SMTP commands are used to contain the envelope information.

### 48.1 Outgoing SMTP and LMTP over TCP/IP

Outgoing SMTP and LMTP over TCP/IP is implemented by the *smtp* transport. The **protocol** option selects which protocol is to be used, but the actual processing is the same in both cases.

If, in response to its EHLO command, Exim is told that the SIZE parameter is supported, it adds SIZE=<*n*> to each subsequent MAIL command. The value of <*n*> is the message size plus the value of the **size\_addition** option (default 1024) to allow for additions to the message such as per-transport header lines, or changes made in a transport filter. If **size\_addition** is set negative, the use of SIZE is suppressed.

If the remote server advertises support for PIPELINING, Exim uses the pipelining extension to SMTP (RFC 2197) to reduce the number of TCP/IP packets required for the transaction.

If the remote server advertises support for the STARTTLS command, and Exim was built to support TLS encryption, it tries to start a TLS session unless the server matches **hosts\_avoid\_tls**. See chapter 42 for more details. Either a match in that or **hosts\_verify\_avoid\_tls** apply when the transport is called for verification.

If the remote server advertises support for the AUTH command, Exim scans the authenticators configuration for any suitable client settings, as described in chapter 33.

Responses from the remote host are supposed to be terminated by CR followed by LF. However, there are known to be hosts that do not send CR characters, so in order to be able to interwork with such hosts, Exim treats LF on its own as a line terminator.

If a message contains a number of different addresses, all those with the same characteristics (for example, the same envelope sender) that resolve to the same set of hosts, in the same order, are sent in a single SMTP transaction, even if they are for different domains, unless there are more than the setting of the **max\_rcpts** option in the *smtp* transport allows, in which case they are split into groups containing no more than **max\_rcpts** addresses each. If **remote\_max\_parallel** is greater than one, such groups may be sent in parallel sessions. The order of hosts with identical MX values is not significant when checking whether addresses can be batched in this way.

When the *smtp* transport suffers a temporary failure that is not message-related, Exim updates its transport-specific database, which contains records indexed by host name that remember which messages are waiting for each particular host. It also updates the retry database with new retry times.

Exim's retry hints are based on host name plus IP address, so if one address of a multi-homed host is broken, it will soon be skipped most of the time. See the next section for more detail about error handling.

When a message is successfully delivered over a TCP/IP SMTP connection, Exim looks in the hints database for the transport to see if there are any queued messages waiting for the host to which it is connected. If it finds one, it creates a new Exim process using the **-MC** option (which can only be used by a process running as root or the Exim user) and passes the TCP/IP socket to it so that it can deliver another message using the same socket. The new process does only those deliveries that are routed to the connected host, and may in turn pass the socket on to a third process, and so on.

The **connection\_max\_messages** option of the *smtp* transport can be used to limit the number of messages sent down a single TCP/IP connection.

The second and subsequent messages delivered down an existing connection are identified in the main log by the addition of an asterisk after the closing square bracket of the IP address.

## 48.2 Errors in outgoing SMTP

Three different kinds of error are recognized for outgoing SMTP: host errors, message errors, and recipient errors.

### Host errors

A host error is not associated with a particular message or with a particular recipient of a message. The host errors are:

- Connection refused or timed out,
- Any error response code on connection,
- Any error response code to EHLO or HELO,
- Loss of connection at any time, except after “.”,
- I/O errors at any time,
- Timeouts during the session, other than in response to MAIL, RCPT or the “.” at the end of the data.

For a host error, a permanent error response on connection, or in response to EHLO, causes all addresses routed to the host to be failed. Any other host error causes all addresses to be deferred, and retry data to be created for the host. It is not tried again, for any message, until its retry time arrives. If the current set of addresses are not all delivered in this run (to some alternative host), the message is added to the list of those waiting for this host, so if it is still undelivered when a subsequent successful delivery is made to the host, it will be sent down the same SMTP connection.

### Message errors

A message error is associated with a particular message when sent to a particular host, but not with a particular recipient of the message. The message errors are:

- Any error response code to MAIL, DATA, or the “.” that terminates the data,
- Timeout after MAIL,
- Timeout or loss of connection after the “.” that terminates the data. A timeout after the DATA command itself is treated as a host error, as is loss of connection at any other time.

For a message error, a permanent error response (5xx) causes all addresses to be failed, and a delivery error report to be returned to the sender. A temporary error response (4xx), or one of the timeouts, causes all addresses to be deferred. Retry data is not created for the host, but instead, a retry record for the combination of host plus message id is created. The message is not added to the list of those waiting for this host. This ensures that the failing message will not be sent to this host again until the retry time arrives. However, other messages that are routed to the host are not affected, so if it is some property of the message that is causing the error, it will not stop the delivery of other mail.

If the remote host specified support for the SIZE parameter in its response to EHLO, Exim adds **SIZE=nnn** to the MAIL command, so an over-large message will cause a message error because the error arrives as a response to MAIL.

## Recipient errors

A recipient error is associated with a particular recipient of a message. The recipient errors are:

- Any error response to RCPT,
- Timeout after RCPT.

For a recipient error, a permanent error response (5xx) causes the recipient address to be failed, and a bounce message to be returned to the sender. A temporary error response (4xx) or a timeout causes the failing address to be deferred, and routing retry data to be created for it. This is used to delay processing of the address in subsequent queue runs, until its routing retry time arrives. This applies to all messages, but because it operates only in queue runs, one attempt will be made to deliver a new message to the failing address before the delay starts to operate. This ensures that, if the failure is really related to the message rather than the recipient (“message too big for this recipient” is a possible example), other messages have a chance of getting delivered. If a delivery to the address does succeed, the retry information gets cleared, so all stuck messages get tried again, and the retry clock is reset.

The message is not added to the list of those waiting for this host. Use of the host for other messages is unaffected, and except in the case of a timeout, other recipients are processed independently, and may be successfully delivered in the current SMTP session. After a timeout it is of course impossible to proceed with the session, so all addresses get deferred. However, those other than the one that failed do not suffer any subsequent retry delays. Therefore, if one recipient is causing trouble, the others have a chance of getting through when a subsequent delivery attempt occurs before the failing recipient’s retry time.

In all cases, if there are other hosts (or IP addresses) available for the current set of addresses (for example, from multiple MX records), they are tried in this run for any undelivered addresses, subject of course to their own retry data. In other words, recipient error retry data does not take effect until the next delivery attempt.

Some hosts have been observed to give temporary error responses to every MAIL command at certain times (“insufficient space” has been seen). It would be nice if such circumstances could be recognized, and defer data for the host itself created, but this is not possible within the current Exim design. What actually happens is that retry data for every (host, message) combination is created.

The reason that timeouts after MAIL and RCPT are treated specially is that these can sometimes arise as a result of the remote host’s verification procedures. Exim makes this assumption, and treats them as if a temporary error response had been received. A timeout after “.” is treated specially because it is known that some broken implementations fail to recognize the end of the message if the last character of the last line is a binary zero. Thus, it is helpful to treat this case as a message error.

Timeouts at other times are treated as host errors, assuming a problem with the host, or the connection to it. If a timeout after MAIL, RCPT, or “.” is really a connection problem, the assumption is that at the next try the timeout is likely to occur at some other point in the dialogue, causing it then to be treated as a host error.

There is experimental evidence that some MTAs drop the connection after the terminating “.” if they do not like the contents of the message for some reason, in contravention of the RFC, which indicates that a 5xx response should be given. That is why Exim treats this case as a message rather than a host error, in order not to delay other messages to the same host.

## 48.3 Incoming SMTP messages over TCP/IP

Incoming SMTP messages can be accepted in one of two ways: by running a listening daemon, or by using *inetd*. In the latter case, the entry in */etc/inetd.conf* should be like this:

```
smtp stream tcp nowait exim /opt/exim/bin/exim in.exim -bs
```

Exim distinguishes between this case and the case of a locally running user agent using the **-bs** option by checking whether or not the standard input is a socket. When it is, either the port must be privileged (less than 1024), or the caller must be root or the Exim user. If any other user passes a



socket with an unprivileged port number, Exim prints a message on the standard error stream and exits with an error code.

By default, Exim does not make a log entry when a remote host connects or disconnects (either via the daemon or *inetd*), unless the disconnection is unexpected. It can be made to write such log entries by setting the **smtp\_connection** log selector.

Commands from the remote host are supposed to be terminated by CR followed by LF. However, there are known to be hosts that do not send CR characters. In order to be able to interwork with such hosts, Exim treats LF on its own as a line terminator. Furthermore, because common code is used for receiving messages from all sources, a CR on its own is also interpreted as a line terminator. However, the sequence “CR, dot, CR” does not terminate incoming SMTP data.

One area that sometimes gives rise to problems concerns the EHLO or HELO commands. Some clients send syntactically invalid versions of these commands, which Exim rejects by default. (This is nothing to do with verifying the data that is sent, so **helo\_verify\_hosts** is not relevant.) You can tell Exim not to apply a syntax check by setting **helo\_accept\_junk\_hosts** to match the broken hosts that send invalid commands.

The amount of disk space available is checked whenever SIZE is received on a MAIL command, independently of whether **message\_size\_limit** or **check\_spool\_space** is configured, unless **smtp\_check\_spool\_space** is set false. A temporary error is given if there is not enough space. If **check\_spool\_space** is set, the check is for that amount of space plus the value given with SIZE, that is, it checks that the addition of the incoming message will not reduce the space below the threshold.

When a message is successfully received, Exim includes the local message id in its response to the final “.” that terminates the data. If the remote host logs this text it can help with tracing what has happened to a message.

The Exim daemon can limit the number of simultaneous incoming connections it is prepared to handle (see the **smtp\_accept\_max** option). It can also limit the number of simultaneous incoming connections from a single remote host (see the **smtp\_accept\_max\_per\_host** option). Additional connection attempts are rejected using the SMTP temporary error code 421.

The Exim daemon does not rely on the SIGCHLD signal to detect when a subprocess has finished, as this can get lost at busy times. Instead, it looks for completed subprocesses every time it wakes up. Provided there are other things happening (new incoming calls, starts of queue runs), completed processes will be noticed and tidied away. On very quiet systems you may sometimes see a “defunct” Exim process hanging about. This is not a problem; it will be noticed when the daemon next wakes up.

When running as a daemon, Exim can reserve some SMTP slots for specific hosts, and can also be set up to reject SMTP calls from non-reserved hosts at times of high system load – for details see the **smtp\_accept\_reserve**, **smtp\_load\_reserve**, and **smtp\_reserve\_hosts** options. The load check applies in both the daemon and *inetd* cases.

Exim normally starts a delivery process for each message received, though this can be varied by means of the **-odq** command line option and the **queue\_only**, **queue\_only\_file**, and **queue\_only\_load** options. The number of simultaneously running delivery processes started in this way from SMTP input can be limited by the **smtp\_accept\_queue** and **smtp\_accept\_queue\_per\_connection** options. When either limit is reached, subsequently received messages are just put on the input queue without starting a delivery process.

The controls that involve counts of incoming SMTP calls (**smtp\_accept\_max**, **smtp\_accept\_queue**, **smtp\_accept\_reserve**) are not available when Exim is started up from the *inetd* daemon, because in that case each connection is handled by an entirely independent Exim process. Control by load average is, however, available with *inetd*.

Exim can be configured to verify addresses in incoming SMTP commands as they are received. See chapter 43 for details. It can also be configured to rewrite addresses at this time – before any syntax checking is done. See section 31.9.

Exim can also be configured to limit the rate at which a client host submits MAIL and RCPT commands in a single SMTP session. See the **smtp\_ratelimit\_hosts** option.

## 48.4 Unrecognized SMTP commands

If Exim receives more than **smtp\_max\_unknown\_commands** unrecognized SMTP commands during a single SMTP connection, it drops the connection after sending the error response to the last command. The default value for **smtp\_max\_unknown\_commands** is 3. This is a defence against some kinds of abuse that subvert web servers into making connections to SMTP ports; in these circumstances, a number of non-SMTP lines are sent first.

## 48.5 Syntax and protocol errors in SMTP commands

A syntax error is detected if an SMTP command is recognized, but there is something syntactically wrong with its data, for example, a malformed email address in a RCPT command. Protocol errors include invalid command sequencing such as RCPT before MAIL. If Exim receives more than **smtp\_max\_synprot\_errors** such commands during a single SMTP connection, it drops the connection after sending the error response to the last command. The default value for **smtp\_max\_synprot\_errors** is 3. This is a defence against broken clients that loop sending bad commands (yes, it has been seen).

## 48.6 Use of non-mail SMTP commands

The “non-mail” SMTP commands are those other than MAIL, RCPT, and DATA. Exim counts such commands, and drops the connection if there are too many of them in a single SMTP session. This action catches some denial-of-service attempts and things like repeated failing AUTHs, or a mad client looping sending EHLO. The global option **smtp\_accept\_max\_nonmail** defines what “too many” means. Its default value is 10.

When a new message is expected, one occurrence of RSET is not counted. This allows a client to send one RSET between messages (this is not necessary, but some clients do it). Exim also allows one uncounted occurrence of HELO or EHLO, and one occurrence of STARTTLS between messages. After starting up a TLS session, another EHLO is expected, and so it too is not counted.

The first occurrence of AUTH in a connection, or immediately following STARTTLS is also not counted. Otherwise, all commands other than MAIL, RCPT, DATA, and QUIT are counted.

You can control which hosts are subject to the limit set by **smtp\_accept\_max\_nonmail** by setting **smtp\_accept\_max\_nonmail\_hosts**. The default value is \*, which makes the limit apply to all hosts. This option means that you can exclude any specific badly-behaved hosts that you have to live with.

## 48.7 The VRFY and EXPN commands

When Exim receives a VRFY or EXPN command on a TCP/IP connection, it runs the ACL specified by **acl\_smtp\_vrfy** or **acl\_smtp\_expn** (as appropriate) in order to decide whether the command should be accepted or not.

When no ACL is defined for VRFY, or if it rejects without setting an explicit response code, the command is accepted (with a 252 SMTP response code) in order to support awkward clients that do a VRFY before every RCPT. When VRFY is accepted, it runs exactly the same code as when Exim is called with the **-bv** option, and returns 250/451/550 SMTP response codes.

If no ACL for EXPN is defined, the command is rejected. When EXPN is accepted, a single-level expansion of the address is done. EXPN is treated as an “address test” (similar to the **-bt** option) rather than a verification (the **-bv** option). If an unqualified local part is given as the argument to EXPN, it is qualified with **qualify\_domain**. Rejections of VRFY and EXPN commands are logged on the main and reject logs, and VRFY verification failures are logged on the main log for consistency with RCPT failures.

## 48.8 The ETRN command

RFC 185 describes an SMTP command called ETRN that is designed to overcome the security problems of the TURN command (which has fallen into disuse). When Exim receives an ETRN command on a TCP/IP connection, it runs the ACL specified by **acl\_smtp\_etrn** in order to decide whether the command should be accepted or not. If no ACL is defined, the command is rejected.

The ETRN command is concerned with “releasing” messages that are awaiting delivery to certain hosts. As Exim does not organize its message queue by host, the only form of ETRN that is supported by default is the one where the text starts with the “#” prefix, in which case the remainder of the text is specific to the SMTP server. A valid ETRN command causes a run of Exim with the **-R** option to happen, with the remainder of the ETRN text as its argument. For example,

```
ETRN #brigadoon
```

runs the command

```
exim -R brigadoon
```

which causes a delivery attempt on all messages with undelivered addresses containing the text “brigadoon”. When **smtp\_etrn\_serialize** is set (the default), Exim prevents the simultaneous execution of more than one queue run for the same argument string as a result of an ETRN command. This stops a misbehaving client from starting more than one queue runner at once.

Exim implements the serialization by means of a hints database in which a record is written whenever a process is started by ETRN, and deleted when the process completes. However, Exim does not keep the SMTP session waiting for the ETRN process to complete. Once ETRN is accepted, the client is sent a “success” return code. Obviously there is scope for hints records to get left lying around if there is a system or program crash. To guard against this, Exim ignores any records that are more than six hours old.

For more control over what ETRN does, the **smtp\_etrn\_command** option can be used. This specifies a command that is run whenever ETRN is received, whatever the form of its argument. For example:

```
smtp_etrn_command = /etc/etrn_command $domain \
                    $sender_host_address
```

The string is split up into arguments which are independently expanded. The expansion variable *\$domain* is set to the argument of the ETRN command, and no syntax checking is done on the contents of this argument. Exim does not wait for the command to complete, so its status code is not checked. Exim runs under its own uid and gid when receiving incoming SMTP, so it is not possible for it to change them before running the command.

## 48.9 Incoming local SMTP

Some user agents use SMTP to pass messages to their local MTA using the standard input and output, as opposed to passing the envelope on the command line and writing the message to the standard input. This is supported by the **-bs** option. This form of SMTP is handled in the same way as incoming messages over TCP/IP (including the use of ACLs), except that the envelope sender given in a MAIL command is ignored unless the caller is trusted. In an ACL you can detect this form of SMTP input by testing for an empty host identification. It is common to have this as the first line in the ACL that runs for RCPT commands:

```
accept hosts = :
```

This accepts SMTP messages from local processes without doing any other tests.

## 48.10 Outgoing batched SMTP

Both the *appendfile* and *pipe* transports can be used for handling batched SMTP. Each has an option called **use\_bsmtplib** which causes messages to be output in BSMTP format. No SMTP responses are possible for this form of delivery. All it is doing is using SMTP commands as a way of transmitting the envelope along with the message.

The message is written to the file or pipe preceded by the SMTP commands MAIL and RCPT, and followed by a line containing a single dot. Lines in the message that start with a dot have an extra dot added. The SMTP command HELO is not normally used. If it is required, the **message\_prefix** option can be used to specify it.

Because *appendfile* and *pipe* are both local transports, they accept only one recipient address at a time by default. However, you can arrange for them to handle several addresses at once by setting the

**batch\_max** option. When this is done for BSMTP, messages may contain multiple RCPT commands. See chapter 25 for more details.

When one or more addresses are routed to a BSMTP transport by a router that sets up a host list, the name of the first host on the list is available to the transport in the variable *\$host*. Here is an example of such a transport and router:

```
begin routers
route_append:
    driver = manualroute
    transport = smtp_appendfile
    route_list = domain.example batch.host.example

begin transports
smtp_appendfile:
    driver = appendfile
    directory = /var/bsmtp/$host
    batch_max = 1000
    use_bsmtp
    user = exim
```

This causes messages addressed to *domain.example* to be written in BSMTP format to */var/bsmtp/batch.host.example*, with only a single copy of each message (unless there are more than 1000 recipients).

## 48.11 Incoming batched SMTP

The **-bs** command line option causes Exim to accept one or more messages by reading SMTP on the standard input, but to generate no responses. If the caller is trusted, the senders in the MAIL commands are believed; otherwise the sender is always the caller of Exim. Unqualified senders and receivers are not rejected (there seems little point) but instead just get qualified. HELO and EHLO act as RSET; VRFY, EXPN, ETRN and HELP, act as NOOP; QUIT quits.

Minimal policy checking is done for BSMTP input. Only the non-SMTP ACL is run in the same way as for non-SMTP local input.

If an error is detected while reading a message, including a missing “.” at the end, Exim gives up immediately. It writes details of the error to the standard output in a stylized way that the calling program should be able to make some use of automatically, for example:

```
554 Unexpected end of file
Transaction started in line 10
Error detected in line 14
```

It writes a more verbose version, for human consumption, to the standard error file, for example:

```
An error was detected while processing a file of BSMTP input.
The error message was:
```

```
501 '>' missing at end of address
```

```
The SMTP transaction started in line 10.
The error was detected in line 12.
The SMTP command at fault was:
```

```
rcpt to:<malformed@in.com.plete
```

```
1 previous message was successfully processed.
The rest of the batch was abandoned.
```

The return code from Exim is zero only if there were no errors. It is 1 if some messages were accepted before an error was detected, and 2 if no messages were accepted.

## 49. Customizing bounce and warning messages

When a message fails to be delivered, or remains on the queue for more than a configured amount of time, Exim sends a message to the original sender, or to an alternative configured address. The text of these messages is built into the code of Exim, but it is possible to change it, either by adding a single string, or by replacing each of the paragraphs by text supplied in a file.

The *From:* and *To:* header lines are automatically generated; you can cause a *Reply-To:* line to be added by setting the **errors\_reply\_to** option. Exim also adds the line

```
Auto-Submitted: auto-generated
```

to all warning and bounce messages,

### 49.1 Customizing bounce messages

If **bounce\_message\_text** is set, its contents are included in the default message immediately after “This message was created automatically by mail delivery software.” The string is not expanded. It is not used if **bounce\_message\_file** is set.

When **bounce\_message\_file** is set, it must point to a template file for constructing error messages. The file consists of a series of text items, separated by lines consisting of exactly four asterisks. If the file cannot be opened, default text is used and a message is written to the main and panic logs. If any text item in the file is empty, default text is used for that item.

Each item of text that is read from the file is expanded, and there are two expansion variables which can be of use here: *\$bounce\_recipient* is set to the recipient of an error message while it is being created, and *\$bounce\_return\_size\_limit* contains the value of the **return\_size\_limit** option, rounded to a whole number.

The items must appear in the file in the following order:

- The first item is included in the headers, and should include at least a *Subject:* header. Exim does not check the syntax of these headers.
- The second item forms the start of the error message. After it, Exim lists the failing addresses with their error messages.
- The third item is used to introduce any text from pipe transports that is to be returned to the sender. It is omitted if there is no such text.
- The fourth, fifth and sixth items will be ignored and may be empty. The fields exist for back-compatibility

The default state (**bounce\_message\_file** unset) is equivalent to the following file, in which the sixth item is empty. The *Subject:* and some other lines have been split in order to fit them on the page:

```
Subject: Mail delivery failed
    ${if eq{$sender_address}{$bounce_recipient}
      {: returning message to sender}}
****
This message was created automatically by mail delivery software.

A message ${if eq{$sender_address}{$bounce_recipient}
  {that you sent }{sent by

<$sender_address>

}}could not be delivered to all of its recipients.
This is a permanent error. The following address(es) failed:
****
The following text was generated during the delivery attempt(s):
****
```

```

----- This is a copy of the message, including all the headers.
-----
****
----- The body of the message is $message_size characters long;
        only the first
----- $bounce_return_size_limit or so are included here.
****

```

## 49.2 Customizing warning messages

The option **warn\_message\_file** can be pointed at a template file for use when warnings about message delays are created. In this case there are only three text sections:

- The first item is included in the headers, and should include at least a *Subject:* header. Exim does not check the syntax of these headers.
- The second item forms the start of the warning message. After it, Exim lists the delayed addresses.
- The third item then ends the message.

The default state is equivalent to the following file, except that some lines have been split here, in order to fit them on the page:

```

Subject: Warning: message $message_exim_id delayed
        $warn_message_delay
****
This message was created automatically by mail delivery software.

A message ${if eq{$sender_address}{$warn_message_recipients}
{that you sent }{sent by

<$sender_address>

}}has not been delivered to all of its recipients after
more than $warn_message_delay on the queue on $primary_hostname.

The message identifier is:      $message_exim_id
The subject of the message is:  $h_subject
The date of the message is:     $h_date

The following address(es) have not yet been delivered:
****
No action is required on your part. Delivery attempts will
continue for some time, and this warning may be repeated at
intervals if the message remains undelivered. Eventually the
mail delivery software will give up, and when that happens,
the message will be returned to you.

```

However, in the default state the subject and date lines are omitted if no appropriate headers exist. During the expansion of this file, *\$warn\_message\_delay* is set to the delay time in one of the forms “<n> minutes” or “<n> hours”, and *\$warn\_message\_recipients* contains a list of recipients for the warning message. There may be more than one if there are multiple addresses with different **errors\_to** settings on the routers that handled them.

## 50. Some common configuration settings

This chapter discusses some configuration settings that seem to be fairly common. More examples and discussion can be found in the Exim book.

### 50.1 Sending mail to a smart host

If you want to send all mail for non-local domains to a “smart host”, you should replace the default *dnslookup* router with a router which does the routing explicitly:

```
send_to_smart_host:
    driver = manualroute
    route_list = !+local_domains smart.host.name
    transport = remote_smtp
```

You can use the smart host’s IP address instead of the name if you wish. If you are using Exim only to submit messages to a smart host, and not for receiving incoming messages, you can arrange for it to do the submission synchronously by setting the **mua\_wrapper** option (see chapter 51).

### 50.2 Using Exim to handle mailing lists

Exim can be used to run simple mailing lists, but for large and/or complicated requirements, the use of additional specialized mailing list software such as Majordomo or Mailman is recommended.

The *redirect* router can be used to handle mailing lists where each list is maintained in a separate file, which can therefore be managed by an independent manager. The **domains** router option can be used to run these lists in a separate domain from normal mail. For example:

```
lists:
    driver = redirect
    domains = lists.example
    file = /usr/lists/$local_part
    forbid_pipe
    forbid_file
    errors_to = $local_part-request@lists.example
    no_more
```

This router is skipped for domains other than *lists.example*. For addresses in that domain, it looks for a file that matches the local part. If there is no such file, the router declines, but because **no\_more** is set, no subsequent routers are tried, and so the whole delivery fails.

The **forbid\_pipe** and **forbid\_file** options prevent a local part from being expanded into a file name or a pipe delivery, which is usually inappropriate in a mailing list.

The **errors\_to** option specifies that any delivery errors caused by addresses taken from a mailing list are to be sent to the given address rather than the original sender of the message. However, before acting on this, Exim verifies the error address, and ignores it if verification fails.

For example, using the configuration above, mail sent to *dicts@lists.example* is passed on to those addresses contained in */usr/lists/dicts*, with error reports directed to *dicts-request@lists.example*, provided that this address can be verified. There could be a file called */usr/lists/dicts-request* containing the address(es) of this particular list’s manager(s), but other approaches, such as setting up an earlier router (possibly using the **local\_part\_prefix** or **local\_part\_suffix** options) to handle addresses of the form **owner-xxx** or **xxx-request**, are also possible.

### 50.3 Syntax errors in mailing lists

If an entry in redirection data contains a syntax error, Exim normally defers delivery of the original address. That means that a syntax error in a mailing list holds up all deliveries to the list. This may not be appropriate when a list is being maintained automatically from data supplied by users, and the addresses are not rigorously checked.

If the **skip\_syntax\_errors** option is set, the *redirect* router just skips entries that fail to parse, noting the incident in the log. If in addition **syntax\_errors\_to** is set to a verifiable address, a message is sent to it whenever a broken address is skipped. It is usually appropriate to set **syntax\_errors\_to** to the same address as **errors\_to**.

## 50.4 Re-expansion of mailing lists

Exim remembers every individual address to which a message has been delivered, in order to avoid duplication, but it normally stores only the original recipient addresses with a message. If all the deliveries to a mailing list cannot be done at the first attempt, the mailing list is re-expanded when the delivery is next tried. This means that alterations to the list are taken into account at each delivery attempt, so addresses that have been added to the list since the message arrived will therefore receive a copy of the message, even though it pre-dates their subscription.

If this behaviour is felt to be undesirable, the **one\_time** option can be set on the *redirect* router. If this is done, any addresses generated by the router that fail to deliver at the first attempt are added to the message as “top level” addresses, and the parent address that generated them is marked “delivered”. Thus, expansion of the mailing list does not happen again at the subsequent delivery attempts. The disadvantage of this is that if any of the failing addresses are incorrect, correcting them in the file has no effect on pre-existing messages.

The original top-level address is remembered with each of the generated addresses, and is output in any log messages. However, any intermediate parent addresses are not recorded. This makes a difference to the log only if the **all\_parents** selector is set, but for mailing lists there is normally only one level of expansion anyway.

## 50.5 Closed mailing lists

The examples so far have assumed open mailing lists, to which anybody may send mail. It is also possible to set up closed lists, where mail is accepted from specified senders only. This is done by making use of the generic **senders** option to restrict the router that handles the list.

The following example uses the same file as a list of recipients and as a list of permitted senders. It requires three routers:

```
lists_request:
  driver = redirect
  domains = lists.example
  local_part_suffix = -request
  file = /usr/lists/$local_part$local_part_suffix
  no_more

lists_post:
  driver = redirect
  domains = lists.example
  senders = ${if exists {/usr/lists/$local_part}\
             {lsearch;/usr/lists/$local_part}{*}}
  file = /usr/lists/$local_part
  forbid_pipe
  forbid_file
  errors_to = $local_part-request@lists.example
  no_more

lists_closed:
  driver = redirect
  domains = lists.example
  allow_fail
  data = :fail: $local_part@lists.example is a closed mailing list
```



All three routers have the same **domains** setting, so for any other domains, they are all skipped. The first router runs only if the local part ends in **-request**. It handles messages to the list manager(s) by means of an open mailing list.

The second router runs only if the **senders** precondition is satisfied. It checks for the existence of a list that corresponds to the local part, and then checks that the sender is on the list by means of a linear search. It is necessary to check for the existence of the file before trying to search it, because otherwise Exim thinks there is a configuration error. If the file does not exist, the expansion of **senders** is **\***, which matches all senders. This means that the router runs, but because there is no list, declines, and **no\_more** ensures that no further routers are run. The address fails with an “unrouteable address” error.

The third router runs only if the second router is skipped, which happens when a mailing list exists, but the sender is not on it. This router forcibly fails the address, giving a suitable error message.

## 50.6 Variable Envelope Return Paths (VERP)

Variable Envelope Return Paths – see <http://cr.yp.to/proto/verp.txt> – are a way of helping mailing list administrators discover which subscription address is the cause of a particular delivery failure. The idea is to encode the original recipient address in the outgoing envelope sender address, so that if the message is forwarded by another host and then subsequently bounces, the original recipient can be extracted from the recipient address of the bounce.

Envelope sender addresses can be modified by Exim using two different facilities: the **errors\_to** option on a router (as shown in previous mailing list examples), or the **return\_path** option on a transport. The second of these is effective only if the message is successfully delivered to another host; it is not used for errors detected on the local host (see the description of **return\_path** in chapter 24). Here is an example of the use of **return\_path** to implement VERP on an *smtp* transport:

```
verp_smtp:
  driver = smtp
  max_rcpt = 1
  return_path = \
    ${if match {$return_path}{^(.+?)~request@your.dom.example\}$}\
      {$1-request+$local_part=$domain@your.dom.example}fail}
```

This has the effect of rewriting the return path (envelope sender) on outgoing SMTP messages, if the local part of the original return path ends in “-request”, and the domain is *your.dom.example*. The rewriting inserts the local part and domain of the recipient into the return path. Suppose, for example, that a message whose return path has been set to *somelist-request@your.dom.example* is sent to *subscriber@other.dom.example*. In the transport, the return path is rewritten as

```
somelist-request+subscriber=other.dom.example@your.dom.example
```

For this to work, you must tell Exim to send multiple copies of messages that have more than one recipient, so that each copy has just one recipient. This is achieved by setting **max\_rcpt** to 1. Without this, a single copy of a message might be sent to several different recipients in the same domain, in which case *\$local\_part* is not available in the transport, because it is not unique.

Unless your host is doing nothing but mailing list deliveries, you should probably use a separate transport for the VERP deliveries, so as not to use extra resources in making one-per-recipient copies for other deliveries. This can easily be done by expanding the **transport** option in the router:

```
dnslookup:
  driver = dnslookup
  domains = ! +local_domains
  transport = \
    ${if match {$return_path}{^(.+?)~request@your.dom.example\}$}\
      {verp_smtp}{remote_smtp}}
  no_more
```

If you want to change the return path using **errors\_to** in a router instead of using **return\_path** in the transport, you need to set **errors\_to** on all routers that handle mailing list addresses. This will ensure that all delivery errors, including those detected on the local host, are sent to the VERP address.

On a host that does no local deliveries and has no manual routing, only the *dnslookup* router needs to be changed. A special transport is not needed for SMTP deliveries. Every mailing list recipient has its own return path value, and so Exim must hand them to the transport one at a time. Here is an example of a *dnslookup* router that implements VERP:

```
verp_dnslookup:
  driver = dnslookup
  domains = ! +local_domains
  transport = remote_smtp
  errors_to = \
    ${if match {$return_path}{^(.+?)-request@your.dom.example\${}}
      {$1-request+$local_part=$domain@your.dom.example} fail}
  no_more
```

Before you start sending out messages with VERPed return paths, you must also configure Exim to accept the bounce messages that come back to those paths. Typically this is done by setting a **local\_part\_suffix** option for a router, and using this to route the messages to wherever you want to handle them.

The overhead incurred in using VERP depends very much on the size of the message, the number of recipient addresses that resolve to the same remote host, and the speed of the connection over which the message is being sent. If a lot of addresses resolve to the same host and the connection is slow, sending a separate copy of the message for each address may take substantially longer than sending a single copy with many recipients (for which VERP cannot be used).

## 50.7 Virtual domains

The phrase *virtual domain* is unfortunately used with two rather different meanings:

- A domain for which there are no real mailboxes; all valid local parts are aliases for other email addresses. Common examples are organizational top-level domains and “vanity” domains.
- One of a number of independent domains that are all handled by the same host, with mailboxes on that host, but where the mailbox owners do not necessarily have login accounts on that host.

The first usage is probably more common, and does seem more “virtual” than the second. This kind of domain can be handled in Exim with a straightforward aliasing router. One approach is to create a separate alias file for each virtual domain. Exim can test for the existence of the alias file to determine whether the domain exists. The *dsearch* lookup type is useful here, leading to a router of this form:

```
virtual:
  driver = redirect
  domains = dsearch;/etc/mail/virtual
  data = ${lookup{$local_part}lsearch{/etc/mail/virtual/$domain}}
  no_more
```

The **domains** option specifies that the router is to be skipped, unless there is a file in the */etc/mail/virtual* directory whose name is the same as the domain that is being processed. When the router runs, it looks up the local part in the file to find a new address (or list of addresses). The **no\_more** setting ensures that if the lookup fails (leading to **data** being an empty string), Exim gives up on the address without trying any subsequent routers.

This one router can handle all the virtual domains because the alias file names follow a fixed pattern. Permissions can be arranged so that appropriate people can edit the different alias files. A successful aliasing operation results in a new envelope recipient address, which is then routed from scratch.

The other kind of “virtual” domain can also be handled in a straightforward way. One approach is to create a file for each domain containing a list of valid local parts, and use it in a router like this:

```

my_domains:
    driver = accept
    domains = dsearch;/etc/mail/domains
    local_parts = lsearch;/etc/mail/domains/$domain
    transport = my_mailboxes

```

The address is accepted if there is a file for the domain, and the local part can be found in the file. The **domains** option is used to check for the file's existence because **domains** is tested before the **local\_parts** option (see section 3.12). You cannot use **require\_files**, because that option is tested after **local\_parts**. The transport is as follows:

```

my_mailboxes:
    driver = appendfile
    file = /var/mail/$domain/$local_part
    user = mail

```

This uses a directory of mailboxes for each domain. The **user** setting is required, to specify which uid is to be used for writing to the mailboxes.

The configuration shown here is just one example of how you might support this requirement. There are many other ways this kind of configuration can be set up, for example, by using a database instead of separate files to hold all the information about the domains.

## 50.8 Multiple user mailboxes

Heavy email users often want to operate with multiple mailboxes, into which incoming mail is automatically sorted. A popular way of handling this is to allow users to use multiple sender addresses, so that replies can easily be identified. Users are permitted to add prefixes or suffixes to their local parts for this purpose. The wildcard facility of the generic router options **local\_part\_prefix** and **local\_part\_suffix** can be used for this. For example, consider this router:

```

userforward:
    driver = redirect
    check_local_user
    file = $home/.forward
    local_part_suffix = -*
    local_part_suffix_optional
    allow_filter

```

It runs a user's *.forward* file for all local parts of the form *username-\**. Within the filter file the user can distinguish different cases by testing the variable *\$local\_part\_suffix*. For example:

```

if $local_part_suffix contains -special then
    save /home/$local_part/Mail/special
endif

```

If the filter file does not exist, or does not deal with such addresses, they fall through to subsequent routers, and, assuming no subsequent use of the **local\_part\_suffix** option is made, they presumably fail. Thus, users have control over which suffixes are valid.

Alternatively, a suffix can be used to trigger the use of a different *.forward* file – which is the way a similar facility is implemented in another MTA:

```

userforward:
    driver = redirect
    check_local_user
    file = $home/.forward$local_part_suffix
    local_part_suffix = -*
    local_part_suffix_optional
    allow_filter

```

If there is no suffix, *forward* is used; if the suffix is *-special*, for example, *forward-special* is used. Once again, if the appropriate file does not exist, or does not deal with the address, it is passed on to subsequent routers, which could, if required, look for an unqualified *forward* file to use as a default.

## 50.9 Simplified vacation processing

The traditional way of running the *vacation* program is for a user to set up a pipe command in a *forward* file (see section 22.6 for syntax details). This is prone to error by inexperienced users. There are two features of Exim that can be used to make this process simpler for users:

- A local part prefix such as “vacation-” can be specified on a router which can cause the message to be delivered directly to the *vacation* program, or alternatively can use Exim’s *autoreply* transport. The contents of a user’s *forward* file are then much simpler. For example:

```
spqr, vacation-spqr
```

- The **require\_files** generic router option can be used to trigger a vacation delivery by checking for the existence of a certain file in the user’s home directory. The **unseen** generic option should also be used, to ensure that the original delivery also proceeds. In this case, all the user has to do is to create a file called, say, *.vacation*, containing a vacation message.

Another advantage of both these methods is that they both work even when the use of arbitrary pipes by users is locked out.

## 50.10 Taking copies of mail

Some installations have policies that require archive copies of all messages to be made. A single copy of each message can easily be taken by an appropriate command in a system filter, which could, for example, use a different file for each day’s messages.

There is also a shadow transport mechanism that can be used to take copies of messages that are successfully delivered by local transports, one copy per delivery. This could be used, *inter alia*, to implement automatic notification of delivery by sites that insist on doing such things.

## 50.11 Intermittently connected hosts

It has become quite common (because it is cheaper) for hosts to connect to the Internet periodically rather than remain connected all the time. The normal arrangement is that mail for such hosts accumulates on a system that is permanently connected.

Exim was designed for use on permanently connected hosts, and so it is not particularly well-suited to use in an intermittently connected environment. Nevertheless there are some features that can be used.

## 50.12 Exim on the upstream server host

It is tempting to arrange for incoming mail for the intermittently connected host to remain on Exim’s queue until the client connects. However, this approach does not scale very well. Two different kinds of waiting message are being mixed up in the same queue – those that cannot be delivered because of some temporary problem, and those that are waiting for their destination host to connect. This makes it hard to manage the queue, as well as wasting resources, because each queue runner scans the entire queue.

A better approach is to separate off those messages that are waiting for an intermittently connected host. This can be done by delivering these messages into local files in batch SMTP, “mailstore”, or other envelope-preserving format, from where they are transmitted by other software when their destination connects. This makes it easy to collect all the mail for one host in a single directory, and to apply local timeout rules on a per-message basis if required.

On a very small scale, leaving the mail on Exim’s queue can be made to work. If you are doing this, you should configure Exim with a long retry period for the intermittent host. For example:

```
cheshire.wonderland.fict.example      *      F, 5d, 24h
```

This stops a lot of failed delivery attempts from occurring, but Exim remembers which messages it has queued up for that host. Once the intermittent host comes online, forcing delivery of one message (either by using the **-M** or **-R** options, or by using the ETRN SMTP command (see section 48.8)) causes all the queued up messages to be delivered, often down a single SMTP connection. While the host remains connected, any new messages get delivered immediately.

If the connecting hosts do not have fixed IP addresses, that is, if a host is issued with a different IP address each time it connects, Exim's retry mechanisms on the holding host get confused, because the IP address is normally used as part of the key string for holding retry information. This can be avoided by unsetting **retry\_include\_ip\_address** on the *smtp* transport. Since this has disadvantages for permanently connected hosts, it is best to arrange a separate transport for the intermittently connected ones.

### 50.13 Exim on the intermittently connected client host

The value of **smtp\_accept\_queue\_per\_connection** should probably be increased, or even set to zero (that is, disabled) on the intermittently connected host, so that all incoming messages down a single connection get delivered immediately.

Mail waiting to be sent from an intermittently connected host will probably not have been routed, because without a connection DNS lookups are not possible. This means that if a normal queue run is done at connection time, each message is likely to be sent in a separate SMTP session. This can be avoided by starting the queue run with a command line option beginning with **-qq** instead of **-q**. In this case, the queue is scanned twice. In the first pass, routing is done but no deliveries take place. The second pass is a normal queue run; since all the messages have been previously routed, those destined for the same host are likely to get sent as multiple deliveries in a single SMTP connection.

## 51. Using Exim as a non-queueing client

On a personal computer, it is a common requirement for all email to be sent to a “smart host”. There are plenty of MUAs that can be configured to operate that way, for all the popular operating systems. However, there are some MUAs for Unix-like systems that cannot be so configured: they submit messages using the command line interface of `/usr/sbin/sendmail`. Furthermore, utility programs such as *cron* submit messages this way.

If the personal computer runs continuously, there is no problem, because it can run a conventional MTA that handles delivery to the smart host, and deal with any delays via its queueing mechanism. However, if the computer does not run continuously or runs different operating systems at different times, queueing email is not desirable.

There is therefore a requirement for something that can provide the `/usr/sbin/sendmail` interface but deliver messages to a smart host without any queueing or retrying facilities. Furthermore, the delivery to the smart host should be synchronous, so that if it fails, the sending MUA is immediately informed. In other words, we want something that extends an MUA that submits to a local MTA via the command line so that it behaves like one that submits to a remote smart host using TCP/SMTP.

There are a number of applications (for example, there is one called *ssmtp*) that do this job. However, people have found them to be lacking in various ways. For instance, you might want to allow aliasing and forwarding to be done before sending a message to the smart host.

Exim already had the necessary infrastructure for doing this job. Just a few tweaks were needed to make it behave as required, though it is somewhat of an overkill to use a fully-featured MTA for this purpose.

There is a Boolean global option called **mua\_wrapper**, defaulting false. Setting **mua\_wrapper** true causes Exim to run in a special mode where it assumes that it is being used to “wrap” a command-line MUA in the manner just described. As well as setting **mua\_wrapper**, you also need to provide a compatible router and transport configuration. Typically there will be just one router and one transport, sending everything to a smart host.

When run in MUA wrapping mode, the behaviour of Exim changes in the following ways:

- A daemon cannot be run, nor will Exim accept incoming messages from *inetd*. In other words, the only way to submit messages is via the command line.
- Each message is synchronously delivered as soon as it is received (**-odi** is assumed). All queueing options (**queue\_only**, **queue\_smtp\_domains**, **control** in an ACL, etc.) are quietly ignored. The Exim reception process does not finish until the delivery attempt is complete. If the delivery is successful, a zero return code is given.
- Address redirection is permitted, but the final routing for all addresses must be to the same remote transport, and to the same list of hosts. Furthermore, the return address (envelope sender) must be the same for all recipients, as must any added or deleted header lines. In other words, it must be possible to deliver the message in a single SMTP transaction, however many recipients there are.
- If these conditions are not met, or if routing any address results in a failure or defer status, or if Exim is unable to deliver all the recipients successfully to one of the smart hosts, delivery of the entire message fails.
- Because no queueing is allowed, all failures are treated as permanent; there is no distinction between 4xx and 5xx SMTP response codes from the smart host. Furthermore, because only a single yes/no response can be given to the caller, it is not possible to deliver to some recipients and not others. If there is an error (temporary or permanent) for any recipient, all are failed.
- If more than one smart host is listed, Exim will try another host after a connection failure or a timeout, in the normal way. However, if this kind of failure happens for all the hosts, the delivery fails.

- When delivery fails, an error message is written to the standard error stream (as well as to Exim's log), and Exim exits to the caller with a return code value 1. The message is expunged from Exim's spool files. No bounce messages are ever generated.
- No retry data is maintained, and any retry rules are ignored.
- A number of Exim options are overridden: **deliver\_drop\_privilege** is forced true, **max\_rcpt** in the *smtp* transport is forced to "unlimited", **remote\_max\_parallel** is forced to one, and fallback hosts are ignored.

The overall effect is that Exim makes a single synchronous attempt to deliver the message, failing if there is any kind of problem. Because no local deliveries are done and no daemon can be run, Exim does not need root privilege. It should be possible to run it *setuid* to *exim* instead of *setuid* to *root*. See section 55.3 for a general discussion about the advantages and disadvantages of running without root privilege.

## 52. Log files

Exim writes three different logs, referred to as the main log, the reject log, and the panic log:

- The main log records the arrival of each message and each delivery in a single line in each case. The format is as compact as possible, in an attempt to keep down the size of log files. Two-character flag sequences make it easy to pick out these lines. A number of other events are recorded in the main log. Some of them are optional, in which case the **log\_selector** option controls whether they are included or not. A Perl script called *eximstats*, which does simple analysis of main log files, is provided in the Exim distribution (see section 53.7).
- The reject log records information from messages that are rejected as a result of a configuration option (that is, for policy reasons). The first line of each rejection is a copy of the line that is also written to the main log. Then, if the message's header has been read at the time the log is written, its contents are written to this log. Only the original header lines are available; header lines added by ACLs are not logged. You can use the reject log to check that your policy controls are working correctly; on a busy host this may be easier than scanning the main log for rejection messages. You can suppress the writing of the reject log by setting **write\_rejectlog** false.
- When certain serious errors occur, Exim writes entries to its panic log. If the error is sufficiently disastrous, Exim bombs out afterwards. Panic log entries are usually written to the main log as well, but can get lost amid the mass of other entries. The panic log should be empty under normal circumstances. It is therefore a good idea to check it (or to have a *cron* script check it) regularly, in order to become aware of any problems. When Exim cannot open its panic log, it tries as a last resort to write to the system log (syslog). This is opened with LOG\_PID+LOG\_CONS and the facility code of LOG\_MAIL. The message itself is written at priority LOG\_CRIT.

Every log line starts with a timestamp, in the format shown in the following example. Note that many of the examples shown in this chapter are line-wrapped. In the log file, this would be all on one line:

```
2001-09-16 16:09:47 SMTP connection from [127.0.0.1] closed
by QUIT
```

By default, the timestamps are in the local timezone. There are two ways of changing this:

- You can set the **timezone** option to a different time zone; in particular, if you set  

```
timezone = UTC
```

the timestamps will be in UTC (aka GMT).
- If you set **log\_timezone** true, the time zone is added to the timestamp, for example:  

```
2003-04-25 11:17:07 +0100 Start queue run: pid=12762
```

Exim does not include its process id in log lines by default, but you can request that it does so by specifying the **pid** log selector (see section 52.15). When this is set, the process id is output, in square brackets, immediately after the time and date.

### 52.1 Where the logs are written

The logs may be written to local files, or to syslog, or both. However, it should be noted that many syslog implementations use UDP as a transport, and are therefore unreliable in the sense that messages are not guaranteed to arrive at the loghost, nor is the ordering of messages necessarily maintained. It has also been reported that on large log files (tens of megabytes) you may need to tweak syslog to prevent it syncing the file with each write – on Linux this has been seen to make syslog take 90% plus of CPU time.

The destination for Exim's logs is configured by setting LOG\_FILE\_PATH in *Local/Makefile* or by setting **log\_file\_path** in the run time configuration. This latter string is expanded, so it can contain, for example, references to the host name:

```
log_file_path = /var/log/$primary_hostname/exim_%slog
```



It is generally advisable, however, to set the string in *Local/Makefile* rather than at run time, because then the setting is available right from the start of Exim's execution. Otherwise, if there's something it wants to log before it has read the configuration file (for example, an error in the configuration file) it will not use the path you want, and may not be able to log at all.

The value of `LOG_FILE_PATH` or **log\_file\_path** is a colon-separated list, currently limited to at most two items. This is one option where the facility for changing a list separator may not be used. The list must always be colon-separated. If an item in the list is "syslog" then syslog is used; otherwise the item must either be an absolute path, containing %s at the point where "main", "reject", or "panic" is to be inserted, or be empty, implying the use of a default path.

When Exim encounters an empty item in the list, it searches the list defined by `LOG_FILE_PATH`, and uses the first item it finds that is neither empty nor "syslog". This means that an empty item in **log\_file\_path** can be used to mean "use the path specified at build time". If no such item exists, log files are written in the *log* subdirectory of the spool directory. This is equivalent to the setting:

```
log_file_path = $spool_directory/log/%slog
```

If you do not specify anything at build time or run time, or if you unset the option at run time (i.e. `log_file_path =` ), that is where the logs are written.

A log file path may also contain %D or %M if datestamped log file names are in use – see section 52.3 below.

Here are some examples of possible settings:

<code>LOG_FILE_PATH=syslog</code>	syslog only
<code>LOG_FILE_PATH=:syslog</code>	syslog and default path
<code>LOG_FILE_PATH=syslog : /usr/log/exim_%s</code>	syslog and specified path
<code>LOG_FILE_PATH=/usr/log/exim_%s</code>	specified path only

If there are more than two paths in the list, the first is used and a panic error is logged.

## 52.2 Logging to local files that are periodically “cycled”

Some operating systems provide centralized and standardized methods for cycling log files. For those that do not, a utility script called *exicyclog* is provided (see section 53.6). This renames and compresses the main and reject logs each time it is called. The maximum number of old logs to keep can be set. It is suggested this script is run as a daily *cron* job.

An Exim delivery process opens the main log when it first needs to write to it, and it keeps the file open in case subsequent entries are required – for example, if a number of different deliveries are being done for the same message. However, remote SMTP deliveries can take a long time, and this means that the file may be kept open long after it is renamed if *exicyclog* or something similar is being used to rename log files on a regular basis. To ensure that a switch of log files is noticed as soon as possible, Exim calls *stat()* on the main log's name before reusing an open file, and if the file does not exist, or its inode has changed, the old file is closed and Exim tries to open the main log from scratch. Thus, an old log file may remain open for quite some time, but no Exim processes should write to it once it has been renamed.

## 52.3 Datestamped log files

Instead of cycling the main and reject log files by renaming them periodically, some sites like to use files whose names contain a timestamp, for example, *mainlog-20031225*. The timestamp is in the form *yyyymmdd* or *yyyymm*. Exim has support for this way of working. It is enabled by setting the **log\_file\_path** option to a path that includes %D or %M at the point where the timestamp is required. For example:

```
log_file_path = /var/spool/exim/log/%slog-%D
log_file_path = /var/log/exim-%s-%D.log
log_file_path = /var/spool/exim/log/%D-%slog
log_file_path = /var/log/exim/%s.%M
```

As before, %s is replaced by “main” or “reject”; the following are examples of names generated by the above examples:

```
/var/spool/exim/log/mainlog-20021225
/var/log/exim-reject-20021225.log
/var/spool/exim/log/20021225-mainlog
/var/log/exim/main.200212
```

When this form of log file is specified, Exim automatically switches to new files at midnight. It does not make any attempt to compress old logs; you will need to write your own script if you require this. You should not run *exicyclog* with this form of logging.

The location of the panic log is also determined by **log\_file\_path**, but it is not dated, because rotation of the panic log does not make sense. When generating the name of the panic log, %D or %M are removed from the string. In addition, if it immediately follows a slash, a following non-alphanumeric character is removed; otherwise a preceding non-alphanumeric character is removed. Thus, the four examples above would give these panic log names:

```
/var/spool/exim/log/paniclog
/var/log/exim-panic.log
/var/spool/exim/log/paniclog
/var/log/exim/panic
```

## 52.4 Logging to syslog

The use of syslog does not change what Exim logs or the format of its messages, except in one respect. If **syslog\_timestamp** is set false, the timestamps on Exim’s log lines are omitted when these lines are sent to syslog. Apart from that, the same strings are written to syslog as to log files. The syslog “facility” is set to LOG\_MAIL, and the program name to “exim” by default, but you can change these by setting the **syslog\_facility** and **syslog\_processname** options, respectively. If Exim was compiled with SYSLOG\_LOG\_PID set in *Local/Makefile* (this is the default in *src/EDITME*), then, on systems that permit it (all except ULTRIX), the LOG\_PID flag is set so that the *syslog()* call adds the pid as well as the time and host name to each line. The three log streams are mapped onto syslog priorities as follows:

- *mainlog* is mapped to LOG\_INFO
- *rejectlog* is mapped to LOG\_NOTICE
- *paniclog* is mapped to LOG\_ALERT

Many log lines are written to both *mainlog* and *rejectlog*, and some are written to both *mainlog* and *paniclog*, so there will be duplicates if these are routed by syslog to the same place. You can suppress this duplication by setting **syslog\_duplication** false.

Exim’s log lines can sometimes be very long, and some of its *rejectlog* entries contain multiple lines when headers are included. To cope with both these cases, entries written to syslog are split into separate *syslog()* calls at each internal newline, and also after a maximum of 870 data characters. (This allows for a total syslog line length of 1024, when additions such as timestamps are added.) If you are running a syslog replacement that can handle lines longer than the 1024 characters allowed by RFC 3164, you should set

```
SYSLOG_LONG_LINES=yes
```

in *Local/Makefile* before building Exim. That stops Exim from splitting long lines, but it still splits at internal newlines in *reject* log entries.

To make it easy to re-assemble split lines later, each component of a split entry starts with a string of the form [*<n>/<m>*] or [*<n>\<m>*] where *<n>* is the component number and *<m>* is the total number of components in the entry. The / delimiter is used when the line was split because it was too long; if it was split because of an internal newline, the \ delimiter is used. For example, supposing the length limit to be 50 instead of 870, the following would be the result of a typical rejection message to *mainlog* (LOG\_INFO), each line in addition being preceded by the time, host name, and pid as added by syslog:

```
[1/5] 2002-09-16 16:09:43 16RdAL-0006pc-00 rejected from
[2/5] [127.0.0.1] (ph10): syntax error in 'From' header
[3/5] when scanning for sender: missing or malformed lo
[4/5] cal part in "<>" (envelope sender is <ph10@cam.exa
[5/5] mple>)
```

The same error might cause the following lines to be written to “rejectlog” (LOG\_NOTICE):

```
[1/18] 2002-09-16 16:09:43 16RdAL-0006pc-00 rejected fro
[2/18] m [127.0.0.1] (ph10): syntax error in 'From' head
[3/18] er when scanning for sender: missing or malformed
[4/18] local part in "<>" (envelope sender is <ph10@cam
[5/18] .example>)
[6/18] Recipients: ph10@some.domain.cam.example
[7/18] P Received: from [127.0.0.1] (ident=ph10)
[8/18] by xxxxx.cam.example with smtp (Exim 4.00)
[9/18] id 16RdAL-0006pc-00
[10/18] for ph10@cam.example; Mon, 16 Sep 2002 16:
[11/18] 09:43 +0100
[12/18] F From: <>
[13/18] Subject: this is a test header
[18/18] X-something: this is another header
[15/18] I Message-Id: <E16RdAL-0006pc-00@xxxxxx.cam.examp
[16/18] le>
[17/18] B Bcc:
[18/18] Date: Mon, 16 Sep 2002 16:09:43 +0100
```

Log lines that are neither too long nor contain newlines are written to syslog without modification.

If only syslog is being used, the Exim monitor is unable to provide a log tail display, unless syslog is routing *mainlog* to a file on the local host and the environment variable EXIMON\_LOG\_FILE\_PATH is set to tell the monitor where it is.

## 52.5 Log line flags

One line is written to the main log for each message received, and for each successful, unsuccessful, and delayed delivery. These lines can readily be picked out by the distinctive two-character flags that immediately follow the timestamp. The flags are:

```
<= message arrival
(= message fakereject
=> normal message delivery
-> additional address in same delivery
>> cutthrough message delivery
*> delivery suppressed by -N
** delivery failed; address bounced
== delivery deferred; temporary problem
```

## 52.6 Logging message reception

The format of the single-line entry in the main log that is written for every message received is shown in the basic example below, which is split over several lines in order to fit it on the page:

```
2002-10-31 08:57:53 16ZCW1-0005MB-00 <= kryten@dwarf.fict.example
H=mailer.fict.example [192.168.123.123] U=exim
P=smtp S=5678 id=<incoming message id>
```

The address immediately following “<=” is the envelope sender address. A bounce message is shown with the sender address “<>”, and if it is locally generated, this is followed by an item of the form

```
R=<message id>
```

which is a reference to the message that caused the bounce to be sent.

For messages from other hosts, the H and U fields identify the remote host and record the RFC 1413 identity of the user that sent the message, if one was received. The number given in square brackets is the IP address of the sending host. If there is a single, unparenthesized host name in the H field, as above, it has been verified to correspond to the IP address (see the **host\_lookup** option). If the name is in parentheses, it was the name quoted by the remote host in the SMTP HELO or EHLO command, and has not been verified. If verification yields a different name to that given for HELO or EHLO, the verified name appears first, followed by the HELO or EHLO name in parentheses.

Misconfigured hosts (and mail forgers) sometimes put an IP address, with or without brackets, in the HELO or EHLO command, leading to entries in the log containing text like these examples:

```
H=(10.21.32.43) [192.168.8.34]
H=([10.21.32.43]) [192.168.8.34]
```

This can be confusing. Only the final address in square brackets can be relied on.

For locally generated messages (that is, messages not received over TCP/IP), the H field is omitted, and the U field contains the login name of the caller of Exim.

For all messages, the P field specifies the protocol used to receive the message. This is the value that is stored in *\$received\_protocol*. In the case of incoming SMTP messages, the value indicates whether or not any SMTP extensions (ESMTP), encryption, or authentication were used. If the SMTP session was encrypted, there is an additional X field that records the cipher suite that was used.

The protocol is set to “esmtpsa” or “esmtpa” for messages received from hosts that have authenticated themselves using the SMTP AUTH command. The first value is used when the SMTP connection was encrypted (“secure”). In this case there is an additional item A= followed by the name of the authenticator that was used. If an authenticated identification was set up by the authenticator’s **server\_set\_id** option, this is logged too, separated by a colon from the authenticator name.

The id field records the existing message id, if present. The size of the received message is given by the S field. When the message is delivered, headers may be removed or added, so that the size of delivered copies of the message may not correspond with this value (and indeed may be different to each other).

The **log\_selector** option can be used to request the logging of additional data when a message is received. See section 52.15 below.

## 52.7 Logging deliveries

The format of the single-line entry in the main log that is written for every delivery is shown in one of the examples below, for local and remote deliveries, respectively. Each example has been split into multiple lines in order to fit it on the page:

```
2002-10-31 08:59:13 16ZCW1-0005MB-00 => marv
<marv@hitch.fict.example> R=localuser T=local_delivery
2002-10-31 09:00:10 16ZCW1-0005MB-00 =>
monk@holistic.fict.example R=dnslookup T=remote_smtp
H=holistic.fict.example [192.168.234.234]
```

For ordinary local deliveries, the original address is given in angle brackets after the final delivery address, which might be a pipe or a file. If intermediate address(es) exist between the original and the final address, the last of these is given in parentheses after the final address. The R and T fields record the router and transport that were used to process the address.

If SMTP AUTH was used for the delivery there is an additional item A= followed by the name of the authenticator that was used. If an authenticated identification was set up by the authenticator’s **client\_set\_id** option, this is logged too, separated by a colon from the authenticator name.

If a shadow transport was run after a successful local delivery, the log line for the successful delivery has an item added on the end, of the form

```
ST=<shadow transport name>
```

If the shadow transport did not succeed, the error message is put in parentheses afterwards.

When more than one address is included in a single delivery (for example, two SMTP RCPT commands in one transaction) the second and subsequent addresses are flagged with -> instead of =>. When two or more messages are delivered down a single SMTP connection, an asterisk follows the IP address in the log lines for the second and subsequent messages.

When two or more messages are delivered down a single TLS connection, the DNS and some TLS-related information logged for the first message delivered will not be present in the log lines for the second and subsequent messages. TLS cipher information is still available.

When delivery is done in cutthrough mode it is flagged with >> and the log line precedes the reception line, since cutthrough waits for a possible rejection from the destination in case it can reject the sourced item.

The generation of a reply message by a filter file gets logged as a “delivery” to the addressee, preceded by “>”.

The **log\_selector** option can be used to request the logging of additional data when a message is delivered. See section 52.15 below.

## 52.8 Discarded deliveries

When a message is discarded as a result of the command “seen finish” being obeyed in a filter file which generates no deliveries, a log entry of the form

```
2002-12-10 00:50:49 16auJc-0001UB-00 => discarded
<low.club@bridge.example> R=userforward
```

is written, to record why no deliveries are logged. When a message is discarded because it is aliased to “:blackhole:” the log line is like this:

```
1999-03-02 09:44:33 10HmaX-0005vi-00 => :blackhole:
<hole@nowhere.example> R=blackhole_router
```

## 52.9 Deferred deliveries

When a delivery is deferred, a line of the following form is logged:

```
2002-12-19 16:20:23 16aiQz-0002Q5-00 == marvin@endrest.example
R=dnslookup T=smtp defer (146): Connection refused
```

In the case of remote deliveries, the error is the one that was given for the last IP address that was tried. Details of individual SMTP failures are also written to the log, so the above line would be preceded by something like

```
2002-12-19 16:20:23 16aiQz-0002Q5-00 Failed to connect to
mail1.endrest.example [192.168.239.239]: Connection refused
```

When a deferred address is skipped because its retry time has not been reached, a message is written to the log, but this can be suppressed by setting an appropriate value in **log\_selector**.

## 52.10 Delivery failures

If a delivery fails because an address cannot be routed, a line of the following form is logged:

```
1995-12-19 16:20:23 0tRiQz-0002Q5-00 ** jim@trek99.example
<jim@trek99.example>: unknown mail domain
```

If a delivery fails at transport time, the router and transport are shown, and the response from the remote host is included, as in this example:

```
2002-07-11 07:14:17 17SXDU-000189-00 ** ace400@pb.example
R=dnslookup T=remote_smtp: SMTP error from remote mailer
after pipelined RCPT TO:<ace400@pb.example>: host
```

```
pbmail3.py.example [192.168.63.111]: 553 5.3.0
<ace400@pb.example>...Addressee unknown
```

The word “pipelined” indicates that the SMTP PIPELINING extension was being used. See **hosts\_****avoid\_esmtp** in the *smtp* transport for a way of disabling PIPELINING. The log lines for all forms of delivery failure are flagged with \*\*.

## 52.11 Fake deliveries

If a delivery does not actually take place because the **-N** option has been used to suppress it, a normal delivery line is written to the log, except that “=>” is replaced by “\*>”.

## 52.12 Completion

A line of the form

```
2002-10-31 09:00:11 16ZCW1-0005MB-00 Completed
```

is written to the main log when a message is about to be removed from the spool at the end of its processing.

## 52.13 Summary of Fields in Log Lines

A summary of the field identifiers that are used in log lines is shown in the following table:

A	authenticator name (and optional id and sender)
C	SMTP confirmation on delivery command list for “no mail in SMTP session”
CV	certificate verification status
D	duration of “no mail in SMTP session”
DN	distinguished name from peer certificate
DS	DNSSEC secured lookups
DT	on => lines: time taken for a delivery
F	sender address (on delivery lines)
H	host name and IP address
I	local interface used
K	CHUNKING extension used
id	message id for incoming message
M8S	8BITMIME status for incoming message
P	on <= lines: protocol used on => and ** lines: return path
PRDR	PRDR extension used
PRX	on <= and => lines: proxy address
Q	alternate queue name
QT	on => lines: time spent on queue so far on “Completed” lines: time spent on queue
R	on <= lines: reference for local bounce on => >> ** and == lines: router name
S	size of message in bytes
SNI	server name indication from TLS client hello
ST	shadow transport name
T	on <= lines: message subject (topic)
TFO	connection took advantage of TCP Fast Open on => ** and == lines: transport name
U	local user or RFC 1413 identity
X	TLS cipher suite

## 52.14 Other log entries

Various other types of log entry are written from time to time. Most should be self-explanatory. Among the more common are:

- *retry time not reached* An address previously suffered a temporary error during routing or local delivery, and the time to retry has not yet arrived. This message is not written to an individual message log file unless it happens during the first delivery attempt.
- *retry time not reached for any host* An address previously suffered temporary errors during remote delivery, and the retry time has not yet arrived for any of the hosts to which it is routed.
- *spool file locked* An attempt to deliver a message cannot proceed because some other Exim process is already working on the message. This can be quite common if queue running processes are started at frequent intervals. The *exiwhat* utility script can be used to find out what Exim processes are doing.
- *error ignored* There are several circumstances that give rise to this message:
  - (1) Exim failed to deliver a bounce message whose age was greater than **ignore\_bounce\_errors\_after**. The bounce was discarded.
  - (2) A filter file set up a delivery using the “noerror” option, and the delivery failed. The delivery was discarded.
  - (3) A delivery set up by a router configured with

```
errors_to = <>
```

failed. The delivery was discarded.

## 52.15 Reducing or increasing what is logged

By setting the **log\_selector** global option, you can disable some of Exim’s default logging, or you can request additional logging. The value of **log\_selector** is made up of names preceded by plus or minus characters. For example:

```
log_selector = +arguments -retry_defer
```

The list of optional log items is in the following table, with the default selection marked by asterisks:

8bitmime	received 8BITMIME status
*acl_warn_skipped	skipped <b>warn</b> statement in ACL
address_rewrite	address rewriting
all_parents	all parents in => lines
arguments	command line arguments
*connection_reject	connection rejections
*delay_delivery	immediate delivery delayed
deliver_time	time taken to perform delivery
delivery_size	add S= <i>nnn</i> to => lines
*dnslist_defer	defers of DNS list (aka RBL) lookups
dnssec	DNSSEC secured lookups
*etrn	ETRN commands
*host_lookup_failed	as it says
ident_timeout	timeout for ident connection
incoming_interface	local interface on <= and => lines
incoming_port	remote port on <= lines
*lost_incoming_connection	as it says (includes timeouts)
millisec	millisecond timestamps and QT,DT,D times
outgoing_interface	local interface on => lines
outgoing_port	add remote port to => lines
*queue_run	start and end queue runs
queue_time	time on queue for one recipient
queue_time_overall	time on queue for whole message

pid	Exim process id
proxy	proxy address on <= and => lines
received_recipients	recipients on <= lines
received_sender	sender on <= lines
*rejected_header	header contents on reject log
*retry_defer	“retry time not reached”
return_path_on_delivery	put return path on => and ** lines
sender_on_delivery	add sender to => lines
*sender_verify_fail	sender verification failures
*size_reject	rejection because too big
*skip_delivery	delivery skipped in a queue run
*smtp_confirmation	SMTP confirmation on => lines
smtp_connection	incoming SMTP connections
smtp_incomplete_transaction	incomplete SMTP transactions
smtp_mailauth	AUTH argument to MAIL commands
smtp_no_mail	session with no MAIL commands
smtp_protocol_error	SMTP protocol errors
smtp_syntax_error	SMTP syntax errors
subject	contents of <i>Subject:</i> on <= lines
*tls_certificate_verified	certificate verification status
*tls_cipher	TLS cipher suite on <= and => lines
tls_peerdn	TLS peer DN on <= and => lines
tls_sni	TLS SNI on <= lines
unknown_in_list	DNS lookup failed in list match
all	all of the above

See also the **slow\_lookup\_log** main configuration option, section 14.4

More details on each of these items follows:

- **8bitmime**: This causes Exim to log any 8BITMIME status of received messages, which may help in tracking down interoperability issues with ancient MTAs that are not 8bit clean. This is added to the “<=” line, tagged with M8S= and a value of 0, 7 or 8, corresponding to “not given”, 7BIT and 8BITMIME respectively.
- **acl\_warn\_skipped**: When an ACL **warn** statement is skipped because one of its conditions cannot be evaluated, a log line to this effect is written if this log selector is set.
- **address\_rewrite**: This applies both to global rewrites and per-transport rewrites, but not to rewrites in filters run as an unprivileged user (because such users cannot access the log).
- **all\_parents**: Normally only the original and final addresses are logged on delivery lines; with this selector, intermediate parents are given in parentheses between them.
- **arguments**: This causes Exim to write the arguments with which it was called to the main log, preceded by the current working directory. This is a debugging feature, added to make it easier to find out how certain MUAs call */usr/sbin/sendmail*. The logging does not happen if Exim has given up root privilege because it was called with the **-C** or **-D** options. Arguments that are empty or that contain white space are quoted. Non-printing characters are shown as escape sequences. This facility cannot log unrecognized arguments, because the arguments are checked before the configuration file is read. The only way to log such cases is to interpose a script such as *util/logargs.sh* between the caller and Exim.
- **connection\_reject**: A log entry is written whenever an incoming SMTP connection is rejected, for whatever reason.
- **delay\_delivery**: A log entry is written whenever a delivery process is not started for an incoming message because the load is too high or too many messages were received on one connection. Logging does not occur if no delivery process is started because **queue\_only** is set or **-odq** was used.



- **deliver\_time**: For each delivery, the amount of real time it has taken to perform the actual delivery is logged as DT=<time>, for example, DT=1s. If millisecond logging is enabled, short times will be shown with greater precision, eg. DT=0.304.
- **delivery\_size**: For each delivery, the size of message delivered is added to the “=>” line, tagged with S=.
- **dnslist\_defer**: A log entry is written if an attempt to look up a host in a DNS black list suffers a temporary error.
- **dnssec**: For message acceptance and (attempted) delivery log lines, when dns lookups gave secure results a tag of DS is added. For acceptance this covers the reverse and forward lookups for host name verification. It does not cover helo-name verification. For delivery this covers the SRV, MX, A and/or AAAA lookups.
- **etrn**: Every valid ETRN command that is received is logged, before the ACL is run to determine whether or not it is actually accepted. An invalid ETRN command, or one received within a message transaction is not logged by this selector (see **smtp\_syntax\_error** and **smtp\_protocol\_error**).
- **host\_lookup\_failed**: When a lookup of a host’s IP addresses fails to find any addresses, or when a lookup of an IP address fails to find a host name, a log line is written. This logging does not apply to direct DNS lookups when routing email addresses, but it does apply to “byname” lookups.
- **ident\_timeout**: A log line is written whenever an attempt to connect to a client’s ident port times out.
- **incoming\_interface**: The interface on which a message was received is added to the “<=” line as an IP address in square brackets, tagged by I= and followed by a colon and the port number. The local interface and port are also added to other SMTP log lines, for example “SMTP connection from”, to rejection lines, and (despite the name) to outgoing “=>” and “->” lines. The latter can be disabled by turning off the **outgoing\_interface** option.
- **proxy**: The internal (closest to the system running Exim) IP address of the proxy, tagged by PRX=, on the “<=” line for a message accepted on a proxied connection or the “=>” line for a message delivered on a proxied connection. See 58.1 for more information.
- **incoming\_port**: The remote port number from which a message was received is added to log entries and *Received*: header lines, following the IP address in square brackets, and separated from it by a colon. This is implemented by changing the value that is put in the *\$sender\_fullhost* and *\$sender\_rcvhost* variables. Recording the remote port number has become more important with the widening use of NAT (see RFC 2505).
- **lost\_incoming\_connection**: A log line is written when an incoming SMTP connection is unexpectedly dropped.
- **millisec**: Timestamps have a period and three decimal places of finer granularity appended to the seconds value.
- **outgoing\_interface**: If **incoming\_interface** is turned on, then the interface on which a message was sent is added to delivery lines as an I= tag followed by IP address in square brackets. You can disable this by turning off the **outgoing\_interface** option.
- **outgoing\_port**: The remote port number is added to delivery log lines (those containing => tags) following the IP address. The local port is also added if **incoming\_interface** and **outgoing\_interface** are both enabled. This option is not included in the default setting, because for most ordinary configurations, the remote port number is always 25 (the SMTP port), and the local port is a random ephemeral port.
- **pid**: The current process id is added to every log line, in square brackets, immediately after the time and date.
- **queue\_run**: The start and end of every queue run are logged.
- **queue\_time**: The amount of time the message has been in the queue on the local host is logged as QT=<time> on delivery (=>) lines, for example, QT=3m45s. The clock starts when Exim starts to

receive the message, so it includes reception time as well as the delivery time for the current address. This means that it may be longer than the difference between the arrival and delivery log line times, because the arrival log line is not written until the message has been successfully received. If millisecond logging is enabled, short times will be shown with greater precision, eg. QT=1.578s.

- **queue\_time\_overall:** The amount of time the message has been in the queue on the local host is logged as QT=<time> on “Completed” lines, for example, QT=3m45s. The clock starts when Exim starts to receive the message, so it includes reception time as well as the total delivery time.
- **received\_recipients:** The recipients of a message are listed in the main log as soon as the message is received. The list appears at the end of the log line that is written when a message is received, preceded by the word “for”. The addresses are listed after they have been qualified, but before any rewriting has taken place. Recipients that were discarded by an ACL for MAIL or RCPT do not appear in the list.
- **received\_sender:** The unrewritten original sender of a message is added to the end of the log line that records the message’s arrival, after the word “from” (before the recipients if **received\_recipients** is also set).
- **rejected\_header:** If a message’s header has been received at the time a rejection is written to the reject log, the complete header is added to the log. Header logging can be turned off individually for messages that are rejected by the *local\_scan()* function (see section 45.2).
- **retry\_defer:** A log line is written if a delivery is deferred because a retry time has not yet been reached. However, this “retry time not reached” message is always omitted from individual message logs after the first delivery attempt.
- **return\_path\_on\_delivery:** The return path that is being transmitted with the message is included in delivery and bounce lines, using the tag P=. This is omitted if no delivery actually happens, for example, if routing fails, or if delivery is to */dev/null* or to *:blackhole:*.
- **sender\_on\_delivery:** The message’s sender address is added to every delivery and bounce line, tagged by F= (for “from”). This is the original sender that was received with the message; it is not necessarily the same as the outgoing return path.
- **sender\_verify\_fail:** If this selector is unset, the separate log line that gives details of a sender verification failure is not written. Log lines for the rejection of SMTP commands contain just “sender verify failed”, so some detail is lost.
- **size\_reject:** A log line is written whenever a message is rejected because it is too big.
- **skip\_delivery:** A log line is written whenever a message is skipped during a queue run because it is frozen or because another process is already delivering it. The message that is written is “spool file is locked”.
- **smtp\_confirmation:** The response to the final “.” in the SMTP or LMTP dialogue for outgoing messages is added to delivery log lines in the form C=<text>. A number of MTAs (including Exim) return an identifying string in this response.
- **smtp\_connection:** A log line is written whenever an incoming SMTP connection is established or closed, unless the connection is from a host that matches **hosts\_connection\_nolog**. (In contrast, **lost\_incoming\_connection** applies only when the closure is unexpected.) This applies to connections from local processes that use **-bs** as well as to TCP/IP connections. If a connection is dropped in the middle of a message, a log line is always written, whether or not this selector is set, but otherwise nothing is written at the start and end of connections unless this selector is enabled.

For TCP/IP connections to an Exim daemon, the current number of connections is included in the log message for each new connection, but note that the count is reset if the daemon is restarted. Also, because connections are closed (and the closure is logged) in subprocesses, the count may not include connections that have been closed but whose termination the daemon has not yet noticed. Thus, while it is possible to match up the opening and closing of connections in the log, the value of the logged counts may not be entirely accurate.

- **smtp\_incomplete\_transaction:** When a mail transaction is aborted by RSET, QUIT, loss of connection, or otherwise, the incident is logged, and the message sender plus any accepted recipients are included in the log line. This can provide evidence of dictionary attacks.
- **smtp\_no\_mail:** A line is written to the main log whenever an accepted SMTP connection terminates without having issued a MAIL command. This includes both the case when the connection is dropped, and the case when QUIT is used. It does not include cases where the connection is rejected right at the start (by an ACL, or because there are too many connections, or whatever). These cases already have their own log lines.

The log line that is written contains the identity of the client in the usual way, followed by D= and a time, which records the duration of the connection. If the connection was authenticated, this fact is logged exactly as it is for an incoming message, with an A= item. If the connection was encrypted, CV=, DN=, and X= items may appear as they do for an incoming message, controlled by the same logging options.

Finally, if any SMTP commands were issued during the connection, a C= item is added to the line, listing the commands that were used. For example,

```
C=EHLO, QUIT
```

shows that the client issued QUIT straight after EHLO. If there were fewer than 20 commands, they are all listed. If there were more than 20 commands, the last 20 are listed, preceded by "...". However, with the default setting of 10 for **smtp\_accept\_max\_nonmail**, the connection will in any case have been aborted before 20 non-mail commands are processed.

- **smtp\_mailauth:** A third subfield with the authenticated sender, colon-separated, is appended to the A= item for a message arrival or delivery log line, if an AUTH argument to the SMTP MAIL command (see 33.2) was accepted or used.
- **smtp\_protocol\_error:** A log line is written for every SMTP protocol error encountered. Exim does not have perfect detection of all protocol errors because of transmission delays and the use of pipelining. If PIPELINING has been advertised to a client, an Exim server assumes that the client will use it, and therefore it does not count “expected” errors (for example, RCPT received after rejecting MAIL) as protocol errors.
- **smtp\_syntax\_error:** A log line is written for every SMTP syntax error encountered. An unrecognized command is treated as a syntax error. For an external connection, the host identity is given; for an internal connection using **-bs** the sender identification (normally the calling user) is given.
- **subject:** The subject of the message is added to the arrival log line, preceded by “T=” (T for “topic”, since S is already used for “size”). Any MIME “words” in the subject are decoded. The **print\_tophitchars** option specifies whether characters with values greater than 127 should be logged unchanged, or whether they should be rendered as escape sequences.
- **tls\_certificate\_verified:** An extra item is added to <= and => log lines when TLS is in use. The item is CV=yes if the peer’s certificate was verified, and CV=no if not.
- **tls\_cipher:** When a message is sent or received over an encrypted connection, the cipher suite used is added to the log line, preceded by X=.
- **tls\_peerdn:** When a message is sent or received over an encrypted connection, and a certificate is supplied by the remote host, the peer DN is added to the log line, preceded by DN=.
- **tls\_sni:** When a message is received over an encrypted connection, and the remote host provided the Server Name Indication extension, the SNI is added to the log line, preceded by SNI=.
- **unknown\_in\_list:** This setting causes a log entry to be written when the result of a list match is failure because a DNS lookup failed.

## 52.16 Message log

In addition to the general log files, Exim writes a log file for each message that it handles. The names of these per-message logs are the message ids, and they are kept in the *msglog* sub-directory of the

spool directory. Each message log contains copies of the log lines that apply to the message. This makes it easier to inspect the status of an individual message without having to search the main log. A message log is deleted when processing of the message is complete, unless **preserve\_message\_logs** is set, but this should be used only with great care because they can fill up your disk very quickly.

On a heavily loaded system, it may be desirable to disable the use of per-message logs, in order to reduce disk I/O. This can be done by setting the **message\_logs** option false.

## 53. Exim utilities

A number of utility scripts and programs are supplied with Exim and are described in this chapter. There is also the Exim Monitor, which is covered in the next chapter. The utilities described here are:

53.1	<i>exiwhat</i>	list what Exim processes are doing
53.2	<i>exiqgrep</i>	grep the queue
53.3	<i>exiqsumm</i>	summarize the queue
53.4	<i>exigrep</i>	search the main log
53.5	<i>exipick</i>	select messages on various criteria
53.6	<i>exicyclog</i>	cycle (rotate) log files
53.7	<i>eximstats</i>	extract statistics from the log
53.8	<i>exim_checkaccess</i>	check address acceptance from given IP
53.9	<i>exim_dbmbuild</i>	build a DBM file
53.10	<i>exinext</i>	extract retry information
53.11	<i>exim_dumpdb</i>	dump a hints database
53.11	<i>exim_tidydb</i>	clean up a hints database
53.11	<i>exim_fixdb</i>	patch a hints database
53.15	<i>exim_lock</i>	lock a mailbox file

Another utility that might be of use to sites with many MTAs is Tom Kistner's *exilog*. It provides log visualizations across multiple Exim servers. See <http://duncanthrax.net/exilog/> for details.

### 53.1 Finding out what Exim processes are doing (exiwhat)

On operating systems that can restart a system call after receiving a signal (most modern OS), an Exim process responds to the SIGUSR1 signal by writing a line describing what it is doing to the file *exim-process.info* in the Exim spool directory. The *exiwhat* script sends the signal to all Exim processes it can find, having first emptied the file. It then waits for one second to allow the Exim processes to react before displaying the results. In order to run *exiwhat* successfully you have to have sufficient privilege to send the signal to the Exim processes, so it is normally run as root.

**Warning:** This is not an efficient process. It is intended for occasional use by system administrators. It is not sensible, for example, to set up a script that sends SIGUSR1 signals to Exim processes at short intervals.

Unfortunately, the *ps* command that *exiwhat* uses to find Exim processes varies in different operating systems. Not only are different options used, but the format of the output is different. For this reason, there are some system configuration options that configure exactly how *exiwhat* works. If it doesn't seem to be working for you, check the following compile-time options:

EXIWHAT_PS_CMD	the command for running <i>ps</i>
EXIWHAT_PS_ARG	the argument for <i>ps</i>
EXIWHAT_EGREP_ARG	the argument for <i>egrep</i> to select from <i>ps</i> output
EXIWHAT_KILL_ARG	the argument for the <i>kill</i> command

An example of typical output from *exiwhat* is

```
164 daemon: -qlh, listening on port 25
10483 running queue: waiting for 0tAycK-0002ij-00 (10492)
10492 delivering 0tAycK-0002ij-00 to mail.ref.example
    [10.19.42.42] (editor@ref.example)
10592 handling incoming call from [192.168.243.242]
10628 accepting a local non-SMTP message
```

The first number in the output line is the process number. The third line has been split here, in order to fit it on the page.

### 53.2 Selective queue listing (exiqgrep)

This utility is a Perl script contributed by Matt Hubbard. It runs

```
exim -bpu
```

or (in case **-a** switch is specified)

```
exim -bp
```

The **-C** option is used to specify an alternate *exim.conf* which might contain alternate exim configuration the queue management might be using.

to obtain a queue listing, and then greps the output to select messages that match given criteria. The following selection options are available:

**-f** <regex>

Match the sender address using a case-insensitive search. The field that is tested is enclosed in angle brackets, so you can test for bounce messages with

```
exiqgrep -f '^<>$'
```

**-r** <regex>

Match a recipient address using a case-insensitive search. The field that is tested is not enclosed in angle brackets.

**-s** <regex>

Match against the size field.

**-y** <seconds>

Match messages that are younger than the given time.

**-o** <seconds>

Match messages that are older than the given time.

**-z**

Match only frozen messages.

**-x**

Match only non-frozen messages.

The following options control the format of the output:

**-c**

Display only the count of matching messages.

**-l**

Long format – display the full message information as output by Exim. This is the default.

**-i**

Display message ids only.

**-b**

Brief format – one line per message.

**-R**

Display messages in reverse order.

**-a**

Include delivered recipients in queue listing.

There is one more option, **-h**, which outputs a list of options.

## 53.3 Summarizing the queue (exiqsumm)

The *exiqsumm* utility is a Perl script which reads the output of `exim -bp` and produces a summary of the messages on the queue. Thus, you use it by running a command such as

```
exim -bp | exiqsumm
```

The output consists of one line for each domain that has messages waiting for it, as in the following example:

Each line lists the number of pending deliveries for a domain, their total volume, and the length of time that the oldest and the newest messages have been waiting. Note that the number of pending deliveries is greater than the number of messages when messages have more than one recipient.

A summary line is output at the end. By default the output is sorted on the domain name, but *exiqsumm* has the options **-a** and **-c**, which cause the output to be sorted by oldest message and by count of messages, respectively. There are also three options that split the messages for each domain into two or more subcounts: **-b** separates bounce messages, **-f** separates frozen messages, and **-s** separates messages according to their sender.

The output of *exim -bp* contains the original addresses in the message, so this also applies to the output from *exiqsumm*. No domains from addresses generated by aliasing or forwarding are included (unless the **one\_time** option of the *redirect* router has been used to convert them into “top level” addresses).

### 53.4 Extracting specific information from the log (exigrep)

The *exigrep* utility is a Perl script that searches one or more main log files for entries that match a given pattern. When it finds a match, it extracts all the log entries for the relevant message, not just those that match the pattern. Thus, *exigrep* can extract complete log entries for a given message, or all mail for a given user, or for a given host, for example. The input files can be in Exim log format or syslog format. If a matching log line is not associated with a specific message, it is included in *exigrep*’s output without any additional lines. The usage is:

```
exigrep [-t<n>] [-I] [-l] [-M] [-v] <pattern> [<log file>] ...
```

If no log file names are given on the command line, the standard input is read.

The **-t** argument specifies a number of seconds. It adds an additional condition for message selection. Messages that are complete are shown only if they spent more than *<n>* seconds on the queue.

By default, *exigrep* does case-insensitive matching. The **-I** option makes it case-sensitive. This may give a performance improvement when searching large log files. Without **-I**, the Perl pattern matches use Perl’s */i* option; with **-I** they do not. In both cases it is possible to change the case sensitivity within the pattern by using *(?i)* or *(?-i)*.

The **-l** option means “literal”, that is, treat all characters in the pattern as standing for themselves. Otherwise the pattern must be a Perl regular expression.

The **-v** option inverts the matching condition. That is, a line is selected if it does *not* match the pattern.

The **-M** options means “related messages”. *exigrep* will show messages that are generated as a result/response to a message that *exigrep* matched normally.

Example of **-M**: user\_a sends a message to user\_b, which generates a bounce back to user\_b. If *exigrep* is used to search for “user\_a”, only the first message will be displayed. But if *exigrep* is used to search for “user\_b”, the first and the second (bounce) message will be displayed. Using **-M** with *exigrep* when searching for “user\_a” will show both messages since the bounce is “related” to or a “result” of the first message that was found by the search term.

If the location of a *zcat* command is known from the definition of *ZCAT\_COMMAND* in *Local/Makefile*, *exigrep* automatically passes any file whose name ends in *COMPRESS\_SUFFIX* through *zcat* as it searches it. If the *ZCAT\_COMMAND* is not executable, *exigrep* tries to use autodetection of some well known compression extensions.

### 53.5 Selecting messages by various criteria (exipick)

John Jetmore’s *exipick* utility is included in the Exim distribution. It lists messages from the queue according to a variety of criteria. For details of *exipick*’s facilities, run *exipick* with the **--help** option.

## 53.6 Cycling log files (exicyclog)

The *exicyclog* script can be used to cycle (rotate) *mainlog* and *rejectlog* files. This is not necessary if only *syslog* is being used, or if you are using log files with timestamps in their names (see section 52.3). Some operating systems have their own standard mechanisms for log cycling, and these can be used instead of *exicyclog* if preferred. There are two command line options for *exicyclog*:

- **-k** <count> specifies the number of log files to keep, overriding the default that is set when Exim is built. The default default is 10.
- **-l** <path> specifies the log file path, in the same format as Exim's **log\_file\_path** option (for example, */var/log/exim\_%slog*), again overriding the script's default, which is to find the setting from Exim's configuration.

Each time *exicyclog* is run the file names get “shuffled down” by one. If the main log file name is *mainlog* (the default) then when *exicyclog* is run *mainlog* becomes *mainlog.01*, the previous *mainlog.01* becomes *mainlog.02* and so on, up to the limit that is set in the script or by the **-k** option. Log files whose numbers exceed the limit are discarded. Reject logs are handled similarly.

If the limit is greater than 99, the script uses 3-digit numbers such as *mainlog.001*, *mainlog.002*, etc. If you change from a number less than 99 to one that is greater, or *vice versa*, you will have to fix the names of any existing log files.

If no *mainlog* file exists, the script does nothing. Files that “drop off” the end are deleted. All files with numbers greater than 01 are compressed, using a compression command which is configured by the **COMPRESS\_COMMAND** setting in *Local/Makefile*. It is usual to run *exicyclog* daily from a root **crontab** entry of the form

```
1 0 * * * su exim -c /usr/exim/bin/exicyclog
```

assuming you have used the name “exim” for the Exim user. You can run *exicyclog* as root if you wish, but there is no need.

## 53.7 Mail statistics (eximstats)

A Perl script called *eximstats* is provided for extracting statistical information from log files. The output is either plain text, or HTML. Exim log files are also supported by the *Lire* system produced by the LogReport Foundation <http://www.logreport.org>.

The *eximstats* script has been hacked about quite a bit over time. The latest version is the result of some extensive revision by Steve Campbell. A lot of information is given by default, but there are options for suppressing various parts of it. Following any options, the arguments to the script are a list of files, which should be main log files. For example:

```
eximstats -nr /var/spool/exim/log/mainlog.01
```

By default, *eximstats* extracts information about the number and volume of messages received from or delivered to various hosts. The information is sorted both by message count and by volume, and the top fifty hosts in each category are listed on the standard output. Similar information, based on email addresses or domains instead of hosts can be requested by means of various options. For messages delivered and received locally, similar statistics are also produced per user.

The output also includes total counts and statistics about delivery errors, and histograms showing the number of messages received and deliveries made in each hour of the day. A delivery with more than one address in its envelope (for example, an SMTP transaction with more than one RCPT command) is counted as a single delivery by *eximstats*.

Though normally more deliveries than receipts are reported (as messages may have multiple recipients), it is possible for *eximstats* to report more messages received than delivered, even though the queue is empty at the start and end of the period in question. If an incoming message contains no valid recipients, no deliveries are recorded for it. A bounce message is handled as an entirely separate message.

*eximstats* always outputs a grand total summary giving the volume and number of messages received and deliveries made, and the number of hosts involved in each case. It also outputs the number of



messages that were delayed (that is, not completely delivered at the first attempt), and the number that had at least one address that failed.

The remainder of the output is in sections that can be independently disabled or modified by various options. It consists of a summary of deliveries by transport, histograms of messages received and delivered per time interval (default per hour), information about the time messages spent on the queue, a list of relayed messages, lists of the top fifty sending hosts, local senders, destination hosts, and destination local users by count and by volume, and a list of delivery errors that occurred.

The relay information lists messages that were actually relayed, that is, they came from a remote host and were directly delivered to some other remote host, without being processed (for example, for aliasing or forwarding) locally.

There are quite a few options for *eximstats* to control exactly what it outputs. These are documented in the Perl script itself, and can be extracted by running the command *perldoc* on the script. For example:

```
perldoc /usr/exim/bin/eximstats
```

## 53.8 Checking access policy (*exim\_checkaccess*)

The **-bh** command line argument allows you to run a fake SMTP session with debugging output, in order to check what Exim is doing when it is applying policy controls to incoming SMTP mail. However, not everybody is sufficiently familiar with the SMTP protocol to be able to make full use of **-bh**, and sometimes you just want to answer the question “Does this address have access?” without bothering with any further details.

The *exim\_checkaccess* utility is a “packaged” version of **-bh**. It takes two arguments, an IP address and an email address:

```
exim_checkaccess 10.9.8.7 A.User@a.domain.example
```

The utility runs a call to Exim with the **-bh** option, to test whether the given email address would be accepted in a RCPT command in a TCP/IP connection from the host with the given IP address. The output of the utility is either the word “accepted”, or the SMTP error response, for example:

```
Rejected:
550 Relay not permitted
```

When running this test, the utility uses <> as the envelope sender address for the MAIL command, but you can change this by providing additional options. These are passed directly to the Exim command. For example, to specify that the test is to be run with the sender address *himself@there.example* you can use:

```
exim_checkaccess 10.9.8.7 A.User@a.domain.example \
-f himself@there.example
```

Note that these additional Exim command line items must be given after the two mandatory arguments.

Because the *exim\_checkaccess* uses **-bh**, it does not perform callouts while running its checks. You can run checks that include callouts by using **-bhc**, but this is not yet available in a “packaged” form.

## 53.9 Making DBM files (*exim\_dbmbuild*)

The *exim\_dbmbuild* program reads an input file containing keys and data in the format used by the *lsearch* lookup (see section 9.3). It writes a DBM file using the lower-cased alias names as keys and the remainder of the information as data. The lower-casing can be prevented by calling the program with the **-nolc** option.

A terminating zero is included as part of the key string. This is expected by the *dbm* lookup type. However, if the option **-nozero** is given, *exim\_dbmbuild* creates files without terminating zeroes in either the key strings or the data strings. The *dbmnz* lookup type can be used with such files.

The program requires two arguments: the name of the input file (which can be a single hyphen to indicate the standard input), and the name of the output file. It creates the output under a temporary name, and then renames it if all went well.

If the native DB interface is in use (`USE_DB` is set in a compile-time configuration file – this is common in free versions of Unix) the two file names must be different, because in this mode the Berkeley DB functions create a single output file using exactly the name given. For example,

```
exim_dbmbuild /etc/aliases /etc/aliases.db
```

reads the system alias file and creates a DBM version of it in */etc/aliases.db*.

In systems that use the *ndbm* routines (mostly proprietary versions of Unix), two files are used, with the suffixes *.dir* and *.pag*. In this environment, the suffixes are added to the second argument of *exim\_dbmbuild*, so it can be the same as the first. This is also the case when the Berkeley functions are used in compatibility mode (though this is not recommended), because in that case it adds a *.db* suffix to the file name.

If a duplicate key is encountered, the program outputs a warning, and when it finishes, its return code is 1 rather than zero, unless the **-noduperr** option is used. By default, only the first of a set of duplicates is used – this makes it compatible with *lsearch* lookups. There is an option **-lastdup** which causes it to use the data for the last duplicate instead. There is also an option **-nowarn**, which stops it listing duplicate keys to **stderr**. For other errors, where it doesn't actually make a new file, the return code is 2.

## 53.10 Finding individual retry times (exinext)

A utility called *exinext* (mostly a Perl script) provides the ability to fish specific information out of the retry database. Given a mail domain (or a complete address), it looks up the hosts for that domain, and outputs any retry information for the hosts or for the domain. At present, the retry information is obtained by running *exim\_dumpdb* (see below) and post-processing the output. For example:

```
$ exinext piglet@milne.fict.example
kanga.milne.example:192.168.8.1 error 146: Connection refused
  first failed: 21-Feb-1996 14:57:34
  last tried:   21-Feb-1996 14:57:34
  next try at:  21-Feb-1996 15:02:34
roo.milne.example:192.168.8.3 error 146: Connection refused
  first failed: 20-Jan-1996 13:12:08
  last tried:   21-Feb-1996 11:42:03
  next try at:  21-Feb-1996 19:42:03
past final cutoff time
```

You can also give *exinext* a local part, without a domain, and it will give any retry information for that local part in your default domain. A message id can be used to obtain retry information pertaining to a specific message. This exists only when an attempt to deliver a message to a remote host suffers a message-specific error (see section 48.2). *exinext* is not particularly efficient, but then it is not expected to be run very often.

The *exinext* utility calls Exim to find out information such as the location of the spool directory. The utility has **-C** and **-D** options, which are passed on to the *exim* commands. The first specifies an alternate Exim configuration file, and the second sets macros for use within the configuration file. These features are mainly to help in testing, but might also be useful in environments where more than one configuration file is in use.

## 53.11 Hints database maintenance

Three utility programs are provided for maintaining the DBM files that Exim uses to contain its delivery hint information. Each program requires two arguments. The first specifies the name of Exim's spool directory, and the second is the name of the database it is to operate on. These are as follows:

- *retry*: the database of retry information
- *wait-<transport name>*: databases of information about messages waiting for remote hosts
- *callout*: the callout cache
- *ratelimit*: the data for implementing the ratelimit ACL condition
- *misc*: other hints data

The *misc* database is used for

- Serializing ETRN runs (when **smtp\_etrn\_serialize** is set)
- Serializing delivery to a specific host (when **serialize\_hosts** is set in an *smtp* transport)
- Limiting the concurrency of specific transports (when **max\_parallel** is set in a transport)

## 53.12 *exim\_dumpdb*

The entire contents of a database are written to the standard output by the *exim\_dumpdb* program, which has no options or arguments other than the spool and database names. For example, to dump the retry database:

```
exim_dumpdb /var/spool/exim retry
```

Two lines of output are produced for each entry:

```
T:mail.ref.example:192.168.242.242 146 77 Connection refused
31-Oct-1995 12:00:12 02-Nov-1995 12:21:39 02-Nov-1995 20:21:39 *
```

The first item on the first line is the key of the record. It starts with one of the letters R, or T, depending on whether it refers to a routing or transport retry. For a local delivery, the next part is the local address; for a remote delivery it is the name of the remote host, followed by its failing IP address (unless **retry\_include\_ip\_address** is set false on the *smtp* transport). If the remote port is not the standard one (port 25), it is added to the IP address. Then there follows an error code, an additional error code, and a textual description of the error.

The three times on the second line are the time of first failure, the time of the last delivery attempt, and the computed time for the next attempt. The line ends with an asterisk if the cutoff time for the last retry rule has been exceeded.

Each output line from *exim\_dumpdb* for the *wait-xxx* databases consists of a host name followed by a list of ids for messages that are or were waiting to be delivered to that host. If there are a very large number for any one host, continuation records, with a sequence number added to the host name, may be seen. The data in these records is often out of date, because a message may be routed to several alternative hosts, and Exim makes no effort to keep cross-references.

## 53.13 *exim\_tidydb*

The *exim\_tidydb* utility program is used to tidy up the contents of a hints database. If run with no options, it removes all records that are more than 30 days old. The age is calculated from the date and time that the record was last updated. Note that, in the case of the retry database, it is *not* the time since the first delivery failure. Information about a host that has been down for more than 30 days will remain in the database, provided that the record is updated sufficiently often.

The cutoff date can be altered by means of the **-t** option, which must be followed by a time. For example, to remove all records older than a week from the retry database:

```
exim_tidydb -t 7d /var/spool/exim retry
```

Both the *wait-xxx* and *retry* databases contain items that involve message ids. In the former these appear as data in records keyed by host – they were messages that were waiting for that host – and in the latter they are the keys for retry information for messages that have suffered certain types of error. When *exim\_tidydb* is run, a check is made to ensure that message ids in database records are those of messages that are still on the queue. Message ids for messages that no longer exist are removed from

*wait-xxx* records, and if this leaves any records empty, they are deleted. For the *retry* database, records whose keys are non-existent message ids are removed. The *exim\_tidydb* utility outputs comments on the standard output whenever it removes information from the database.

Certain records are automatically removed by Exim when they are no longer needed, but others are not. For example, if all the MX hosts for a domain are down, a retry record is created for each one. If the primary MX host comes back first, its record is removed when Exim successfully delivers to it, but the records for the others remain because Exim has not tried to use those hosts.

It is important, therefore, to run *exim\_tidydb* periodically on all the hints databases. You should do this at a quiet time of day, because it requires a database to be locked (and therefore inaccessible to Exim) while it does its work. Removing records from a DBM file does not normally make the file smaller, but all the common DBM libraries are able to re-use the space that is released. After an initial phase of increasing in size, the databases normally reach a point at which they no longer get any bigger, as long as they are regularly tidied.

**Warning:** If you never run *exim\_tidydb*, the space used by the hints databases is likely to keep on increasing.

## 53.14 *exim\_fixdb*

The *exim\_fixdb* program is a utility for interactively modifying databases. Its main use is for testing Exim, but it might also be occasionally useful for getting round problems in a live system. It has no options, and its interface is somewhat crude. On entry, it prompts for input with a right angle-bracket. A key of a database record can then be entered, and the data for that record is displayed.

If “d” is typed at the next prompt, the entire record is deleted. For all except the *retry* database, that is the only operation that can be carried out. For the *retry* database, each field is output preceded by a number, and data for individual fields can be changed by typing the field number followed by new data, for example:

```
> 4 951102:1000
```

resets the time of the next delivery attempt. Time values are given as a sequence of digit pairs for year, month, day, hour, and minute. Colons can be used as optional separators.

## 53.15 Mailbox maintenance (*exim\_lock*)

The *exim\_lock* utility locks a mailbox file using the same algorithm as Exim. For a discussion of locking issues, see section 26.3. *Exim\_lock* can be used to prevent any modification of a mailbox by Exim or a user agent while investigating a problem. The utility requires the name of the file as its first argument. If the locking is successful, the second argument is run as a command (using C’s *system()* function); if there is no second argument, the value of the SHELL environment variable is used; if this is unset or empty, */bin/sh* is run. When the command finishes, the mailbox is unlocked and the utility ends. The following options are available:

### **-fcntl**

Use *fcntl()* locking on the open mailbox.

### **-flock**

Use *flock()* locking on the open mailbox, provided the operating system supports it.

### **-interval**

This must be followed by a number, which is a number of seconds; it sets the interval to sleep between retries (default 3).

### **-lockfile**

Create a lock file before opening the mailbox.

### **-mbx**

Lock the mailbox using MBX rules.

### **-q**

Suppress verification output.

**-retries**

This must be followed by a number; it sets the number of times to try to get the lock (default 10).

**-restore\_time**

This option causes **exim\_lock** to restore the modified and read times to the locked file before exiting. This allows you to access a locked mailbox (for example, to take a backup copy) without disturbing the times that the user subsequently sees.

**-timeout**

This must be followed by a number, which is a number of seconds; it sets a timeout to be used with a blocking *fcntl()* lock. If it is not set (the default), a non-blocking call is used.

**-v**

Generate verbose output.

If none of **-fcntl**, **-flock**, **-lockfile** or **-mbx** are given, the default is to create a lock file and also to use *fcntl()* locking on the mailbox, which is the same as Exim's default. The use of **-flock** or **-fcntl** requires that the file be writeable; the use of **-lockfile** requires that the directory containing the file be writeable. Locking by lock file does not last for ever; Exim assumes that a lock file is expired if it is more than 30 minutes old.

The **-mbx** option can be used with either or both of **-fcntl** or **-flock**. It assumes **-fcntl** by default. MBX locking causes a shared lock to be taken out on the open mailbox, and an exclusive lock on the file */tmp/.n.m* where *n* and *m* are the device number and inode number of the mailbox file. When the locking is released, if an exclusive lock can be obtained for the mailbox, the file in */tmp* is deleted.

The default output contains verification of the locking that takes place. The **-v** option causes some additional information to be given. The **-q** option suppresses all output except error messages.

A command such as

```
exim_lock /var/spool/mail/spqr
```

runs an interactive shell while the file is locked, whereas

```
exim_lock -q /var/spool/mail/spqr <<End
<some commands>
End
```

runs a specific non-interactive sequence of commands while the file is locked, suppressing all verification output. A single command can be run by a command such as

```
exim_lock -q /var/spool/mail/spqr \
    "cp /var/spool/mail/spqr /some/where"
```

Note that if a command is supplied, it must be entirely contained within the second argument – hence the quotes.

## 54. The Exim monitor

The Exim monitor is an application which displays in an X window information about the state of Exim's queue and what Exim is doing. An admin user can perform certain operations on messages from this GUI interface; however all such facilities are also available from the command line, and indeed, the monitor itself makes use of the command line to perform any actions requested.

### 54.1 Running the monitor

The monitor is started by running the script called *eximon*. This is a shell script that sets up a number of environment variables, and then runs the binary called *eximon.bin*. The default appearance of the monitor window can be changed by editing the *Local/eximon.conf* file created by editing *exim\_monitor/EDITME*. Comments in that file describe what the various parameters are for.

The parameters that get built into the *eximon* script can be overridden for a particular invocation by setting up environment variables of the same names, preceded by EXIMON\_. For example, a shell command such as

```
EXIMON_LOG_DEPTH=400 eximon
```

(in a Bourne-compatible shell) runs *eximon* with an overriding setting of the LOG\_DEPTH parameter. If EXIMON\_LOG\_FILE\_PATH is set in the environment, it overrides the Exim log file configuration. This makes it possible to have *eximon* tailing log data that is written to syslog, provided that MAIL.INFO syslog messages are routed to a file on the local host.

X resources can be used to change the appearance of the window in the normal way. For example, a resource setting of the form

```
Eximon*background: gray94
```

changes the colour of the background to light grey rather than white. The stripcharts are drawn with both the data lines and the reference lines in black. This means that the reference lines are not visible when on top of the data. However, their colour can be changed by setting a resource called "highlight" (an odd name, but that's what the Athena stripchart widget uses). For example, if your X server is running Unix, you could set up lighter reference lines in the stripcharts by obeying

```
xrdb -merge <<End
Eximon*highlight: gray
End
```

In order to see the contents of messages on the queue, and to operate on them, *eximon* must either be run as root or by an admin user.

The command-line parameters of *eximon* are passed to *eximon.bin* and may contain X11 resource parameters interpreted by the X11 library. In addition, if the first parameter starts with the string "gdb" then it is removed and the binary is invoked under gdb (the parameter is used as the gdb command-name, so versioned variants of gdb can be invoked).

The monitor's window is divided into three parts. The first contains one or more stripcharts and two action buttons, the second contains a "tail" of the main log file, and the third is a display of the queue of messages awaiting delivery, with two more action buttons. The following sections describe these different parts of the display.

### 54.2 The stripcharts

The first stripchart is always a count of messages on the queue. Its name can be configured by setting QUEUE\_STRIPCHART\_NAME in the *Local/eximon.conf* file. The remaining stripcharts are defined in the configuration script by regular expression matches on log file entries, making it possible to display, for example, counts of messages delivered to certain hosts or using certain transports. The supplied defaults display counts of received and delivered messages, and of local and SMTP deliveries. The default period between stripchart updates is one minute; this can be adjusted by a parameter in the *Local/eximon.conf* file.

The stripchart displays rescale themselves automatically as the value they are displaying changes. There are always 10 horizontal lines in each chart; the title string indicates the value of each division when it is greater than one. For example, “x2” means that each division represents a value of 2.

It is also possible to have a stripchart which shows the percentage fullness of a particular disk partition, which is useful when local deliveries are confined to a single partition.

This relies on the availability of the *statvfs()* function or equivalent in the operating system. Most, but not all versions of Unix that support Exim have this. For this particular stripchart, the top of the chart always represents 100%, and the scale is given as “x10%”. This chart is configured by setting *SIZE\_STRIPCHART* and (optionally) *SIZE\_STRIPCHART\_NAME* in the *Local/eximon.conf* file.

## 54.3 Main action buttons

Below the stripcharts there is an action button for quitting the monitor. Next to this is another button marked “Size”. They are placed here so that shrinking the window to its default minimum size leaves just the queue count stripchart and these two buttons visible. Pressing the “Size” button causes the window to expand to its maximum size, unless it is already at the maximum, in which case it is reduced to its minimum.

When expanding to the maximum, if the window cannot be fully seen where it currently is, it is moved back to where it was the last time it was at full size. When it is expanding from its minimum size, the old position is remembered, and next time it is reduced to the minimum it is moved back there.

The idea is that you can keep a reduced window just showing one or two stripcharts at a convenient place on your screen, easily expand it to show the full window when required, and just as easily put it back to what it was. The idea is copied from what the *twm* window manager does for its *f.fullzoom* action. The minimum size of the window can be changed by setting the *MIN\_HEIGHT* and *MIN\_WIDTH* values in *Local/eximon.conf*.

Normally, the monitor starts up with the window at its full size, but it can be built so that it starts up with the window at its smallest size, by setting *START\_SMALL=yes* in *Local/eximon.conf*.

## 54.4 The log display

The second section of the window is an area in which a display of the tail of the main log is maintained. To save space on the screen, the timestamp on each log line is shortened by removing the date and, if *log\_timezone* is set, the timezone. The log tail is not available when the only destination for logging data is syslog, unless the syslog lines are routed to a local file whose name is passed to *eximon* via the *EXIMON\_LOG\_FILE\_PATH* environment variable.

The log sub-window has a scroll bar at its lefthand side which can be used to move back to look at earlier text, and the up and down arrow keys also have a scrolling effect. The amount of log that is kept depends on the setting of *LOG\_BUFFER* in *Local/eximon.conf*, which specifies the amount of memory to use. When this is full, the earlier 50% of data is discarded – this is much more efficient than throwing it away line by line. The sub-window also has a horizontal scroll bar for accessing the ends of long log lines. This is the only means of horizontal scrolling; the right and left arrow keys are not available. Text can be cut from this part of the window using the mouse in the normal way. The size of this subwindow is controlled by parameters in the configuration file *Local/eximon.conf*.

Searches of the text in the log window can be carried out by means of the ^R and ^S keystrokes, which default to a reverse and a forward search, respectively. The search covers only the text that is displayed in the window. It cannot go further back up the log.

The point from which the search starts is indicated by a caret marker. This is normally at the end of the text in the window, but can be positioned explicitly by pointing and clicking with the left mouse button, and is moved automatically by a successful search. If new text arrives in the window when it is scrolled back, the caret remains where it is, but if the window is not scrolled back, the caret is moved to the end of the new text.

Pressing ^R or ^S pops up a window into which the search text can be typed. There are buttons for selecting forward or reverse searching, for carrying out the search, and for cancelling. If the “Search” button is pressed, the search happens and the window remains so that further searches can be done. If the “Return” key is pressed, a single search is done and the window is closed. If ^C is typed the search is cancelled.

The searching facility is implemented using the facilities of the Athena text widget. By default this pops up a window containing both “search” and “replace” options. In order to suppress the unwanted “replace” portion for eximon, a modified version of the **TextPop** widget is distributed with Exim. However, the linkers in BSDI and HP-UX seem unable to handle an externally provided version of **TextPop** when the remaining parts of the text widget come from the standard libraries. The compile-time option EXIMON\_TEXTPOP can be unset to cut out the modified **TextPop**, making it possible to build Eximon on these systems, at the expense of having unwanted items in the search popup window.

## 54.5 The queue display

The bottom section of the monitor window contains a list of all messages that are on the queue, which includes those currently being received or delivered, as well as those awaiting delivery. The size of this subwindow is controlled by parameters in the configuration file *Local/eximon.conf*, and the frequency at which it is updated is controlled by another parameter in the same file – the default is 5 minutes, since queue scans can be quite expensive. However, there is an “Update” action button just above the display which can be used to force an update of the queue display at any time.

When a host is down for some time, a lot of pending mail can build up for it, and this can make it hard to deal with other messages on the queue. To help with this situation there is a button next to “Update” called “Hide”. If pressed, a dialogue box called “Hide addresses ending with” is put up. If you type anything in here and press “Return”, the text is added to a chain of such texts, and if every undelivered address in a message matches at least one of the texts, the message is not displayed.

If there is an address that does not match any of the texts, all the addresses are displayed as normal. The matching happens on the ends of addresses so, for example, *cam.ac.uk* specifies all addresses in Cambridge, while *xxx@foo.com.example* specifies just one specific address. When any hiding has been set up, a button called “Unhide” is displayed. If pressed, it cancels all hiding. Also, to ensure that hidden messages do not get forgotten, a hide request is automatically cancelled after one hour.

While the dialogue box is displayed, you can’t press any buttons or do anything else to the monitor window. For this reason, if you want to cut text from the queue display to use in the dialogue box, you have to do the cutting before pressing the “Hide” button.

The queue display contains, for each unhidden queued message, the length of time it has been on the queue, the size of the message, the message id, the message sender, and the first undelivered recipient, all on one line. If it is a bounce message, the sender is shown as “<>”. If there is more than one recipient to which the message has not yet been delivered, subsequent ones are listed on additional lines, up to a maximum configured number, following which an ellipsis is displayed. Recipients that have already received the message are not shown.

If a message is frozen, an asterisk is displayed at the left-hand side.

The queue display has a vertical scroll bar, and can also be scrolled by means of the arrow keys. Text can be cut from it using the mouse in the normal way. The text searching facilities, as described above for the log window, are also available, but the caret is always moved to the end of the text when the queue display is updated.

## 54.6 The queue menu

If the **shift** key is held down and the left button is clicked when the mouse pointer is over the text for any message, an action menu pops up, and the first line of the queue display for the message is highlighted. This does not affect any selected text.

If you want to use some other event for popping up the menu, you can set the MENU\_EVENT parameter in *Local/eximon.conf* to change the default, or set EXIMON\_MENU\_EVENT in the



environment before starting the monitor. The value set in this parameter is a standard X event description. For example, to run `eximon` using **ctrl** rather than **shift** you could use

```
EXIMON_MENU_EVENT='Ctrl<Btn1Down>' eximon
```

The title of the menu is the message id, and it contains entries which act as follows:

- *message log*: The contents of the message log for the message are displayed in a new text window.
- *headers*: Information from the spool file that contains the envelope information and headers is displayed in a new text window. See chapter 56 for a description of the format of spool files.
- *body*: The contents of the spool file containing the body of the message are displayed in a new text window. There is a default limit of 20,000 bytes to the amount of data displayed. This can be changed by setting the `BODY_MAX` option at compile time, or the `EXIMON_BODY_MAX` option at run time.
- *deliver message*: A call to Exim is made using the **-M** option to request delivery of the message. This causes an automatic thaw if the message is frozen. The **-v** option is also set, and the output from Exim is displayed in a new text window. The delivery is run in a separate process, to avoid holding up the monitor while the delivery proceeds.
- *freeze message*: A call to Exim is made using the **-Mf** option to request that the message be frozen.
- *thaw message*: A call to Exim is made using the **-Mt** option to request that the message be thawed.
- *give up on msg*: A call to Exim is made using the **-Mg** option to request that Exim gives up trying to deliver the message. A bounce message is generated for any remaining undelivered addresses.
- *remove message*: A call to Exim is made using the **-Mrm** option to request that the message be deleted from the system without generating a bounce message.
- *add recipient*: A dialog box is displayed into which a recipient address can be typed. If the address is not qualified and the `QUALIFY_DOMAIN` parameter is set in *Local/eximon.conf*, the address is qualified with that domain. Otherwise it must be entered as a fully qualified address. Pressing RETURN causes a call to Exim to be made using the **-Mar** option to request that an additional recipient be added to the message, unless the entry box is empty, in which case no action is taken.
- *mark delivered*: A dialog box is displayed into which a recipient address can be typed. If the address is not qualified and the `QUALIFY_DOMAIN` parameter is set in *Local/eximon.conf*, the address is qualified with that domain. Otherwise it must be entered as a fully qualified address. Pressing RETURN causes a call to Exim to be made using the **-Mmd** option to mark the given recipient address as already delivered, unless the entry box is empty, in which case no action is taken.
- *mark all delivered*: A call to Exim is made using the **-Mmad** option to mark all recipient addresses as already delivered.
- *edit sender*: A dialog box is displayed initialized with the current sender's address. Pressing RETURN causes a call to Exim to be made using the **-Mes** option to replace the sender address, unless the entry box is empty, in which case no action is taken. If you want to set an empty sender (as in bounce messages), you must specify it as "<>". Otherwise, if the address is not qualified and the `QUALIFY_DOMAIN` parameter is set in *Local/eximon.conf*, the address is qualified with that domain.

When a delivery is forced, a window showing the **-v** output is displayed. In other cases when a call to Exim is made, if there is any output from Exim (in particular, if the command fails) a window containing the command and the output is displayed. Otherwise, the results of the action are normally apparent from the log and queue displays. However, if you set `ACTION_OUTPUT=yes` in *Local/eximon.conf*, a window showing the Exim command is always opened, even if no output is generated.

The queue display is automatically updated for actions such as freezing and thawing, unless `ACTION_QUEUE_UPDATE=no` has been set in *Local/eximon.conf*. In this case the "Update" button has to be used to force an update of the display after one of these actions.

In any text window that is displayed as result of a menu action, the normal cut-and-paste facility is available, and searching can be carried out using ^R and ^S, as described above for the log tail window.

## 55. Security considerations

This chapter discusses a number of issues concerned with security, some of which are also covered in other parts of this manual.

For reasons that this author does not understand, some people have promoted Exim as a “particularly secure” mailer. Perhaps it is because of the existence of this chapter in the documentation. However, the intent of the chapter is simply to describe the way Exim works in relation to certain security concerns, not to make any specific claims about the effectiveness of its security as compared with other MTAs.

What follows is a description of the way Exim is supposed to be. Best efforts have been made to try to ensure that the code agrees with the theory, but an absence of bugs can never be guaranteed. Any that are reported will get fixed as soon as possible.

### 55.1 Building a more “hardened” Exim

There are a number of build-time options that can be set in *Local/Makefile* to create Exim binaries that are “harder” to attack, in particular by a rogue Exim administrator who does not have the root password, or by someone who has penetrated the Exim (but not the root) account. These options are as follows:

- `ALT_CONFIG_PREFIX` can be set to a string that is required to match the start of any file names used with the `-C` option. When it is set, these file names are also not allowed to contain the sequence `“/./”`. (However, if the value of the `-C` option is identical to the value of `CONFIGURE_FILE` in *Local/Makefile*, Exim ignores `-C` and proceeds as usual.) There is no default setting for `ALT_CONFIG_PREFIX`.

If the permitted configuration files are confined to a directory to which only root has access, this guards against someone who has broken into the Exim account from running a privileged Exim with an arbitrary configuration file, and using it to break into other accounts.

- If a non-trusted configuration file (i.e. not the default configuration file or one which is trusted by virtue of being listed in the `TRUSTED_CONFIG_LIST` file) is specified with `-C`, or if macros are given with `-D` (but see the next item), then root privilege is retained only if the caller of Exim is root. This locks out the possibility of testing a configuration using `-C` right through message reception and delivery, even if the caller is root. The reception works, but by that time, Exim is running as the Exim user, so when it re-execs to regain privilege for the delivery, the use of `-C` causes privilege to be lost. However, root can test reception and delivery using two separate commands.
- The `WHITELIST_D_MACROS` build option declares some macros to be safe to override with `-D` if the real uid is one of root, the Exim run-time user or the `CONFIGURE_OWNER`, if defined. The potential impact of this option is limited by requiring the run-time value supplied to `-D` to match a regex that errs on the restrictive side. Requiring build-time selection of safe macros is onerous but this option is intended solely as a transition mechanism to permit previously-working configurations to continue to work after release 4.73.
- If `DISABLE_D_OPTION` is defined, the use of the `-D` command line option is disabled.
- `FIXED_NEVER_USERS` can be set to a colon-separated list of users that are never to be used for any deliveries. This is like the `never_users` runtime option, but it cannot be overridden; the runtime option adds additional users to the list. The default setting is `“root”`; this prevents a non-root user who is permitted to modify the runtime file from using Exim as a way to get root.

### 55.2 Root privilege

The Exim binary is normally setuid to root, which means that it gains root privilege (runs as root) when it starts execution. In some special cases (for example, when the daemon is not in use and there are no local deliveries), it may be possible to run Exim setuid to some user other than root. This is discussed in the next section. However, in most installations, root privilege is required for two things:

- To set up a socket connected to the standard SMTP port (25) when initialising the listening daemon. If Exim is run from *inetd*, this privileged action is not required.
- To be able to change uid and gid in order to read users' *.forward* files and perform local deliveries as the receiving user or as specified in the configuration.

It is not necessary to be root to do any of the other things Exim does, such as receiving messages and delivering them externally over SMTP, and it is obviously more secure if Exim does not run as root except when necessary. For this reason, a user and group for Exim to use must be defined in *Local/Makefile*. These are known as “the Exim user” and “the Exim group”. Their values can be changed by the run time configuration, though this is not recommended. Often a user called *exim* is used, but some sites use *mail* or another user name altogether.

Exim uses *setuid()* whenever it gives up root privilege. This is a permanent abdication; the process cannot regain root afterwards. Prior to release 4.00, *seteuid()* was used in some circumstances, but this is no longer the case.

After a new Exim process has interpreted its command line options, it changes uid and gid in the following cases:

- If the **-C** option is used to specify an alternate configuration file, or if the **-D** option is used to define macro values for the configuration, and the calling process is not running as root, the uid and gid are changed to those of the calling process. However, if `DISABLE_D_OPTION` is defined in *Local/Makefile*, the **-D** option may not be used at all. If `WHITELIST_D_MACROS` is defined in *Local/Makefile*, then some macro values can be supplied if the calling process is running as root, the Exim run-time user or `CONFIGURE_OWNER`, if defined.
- If the expansion test option (**-be**) or one of the filter testing options (**-bf** or **-bF**) are used, the uid and gid are changed to those of the calling process.
- If the process is not a daemon process or a queue runner process or a delivery process or a process for testing address routing (started with **-bt**), the uid and gid are changed to the Exim user and group. This means that Exim always runs under its own uid and gid when receiving messages. This also applies when testing address verification (the **-bv** option) and testing incoming message policy controls (the **-bh** option).
- For a daemon, queue runner, delivery, or address testing process, the uid remains as root at this stage, but the gid is changed to the Exim group.

The processes that initially retain root privilege behave as follows:

- A daemon process changes the gid to the Exim group and the uid to the Exim user after setting up one or more listening sockets. The *initgroups()* function is called, so that if the Exim user is in any additional groups, they will be used during message reception.
- A queue runner process retains root privilege throughout its execution. Its job is to fork a controlled sequence of delivery processes.
- A delivery process retains root privilege throughout most of its execution, but any actual deliveries (that is, the transports themselves) are run in subprocesses which always change to a non-root uid and gid. For local deliveries this is typically the uid and gid of the owner of the mailbox; for remote deliveries, the Exim uid and gid are used. Once all the delivery subprocesses have been run, a delivery process changes to the Exim uid and gid while doing post-delivery tidying up such as updating the retry database and generating bounce and warning messages.

While the recipient addresses in a message are being routed, the delivery process runs as root. However, if a user's filter file has to be processed, this is done in a subprocess that runs under the individual user's uid and gid. A system filter is run as root unless `system_filter_user` is set.

- A process that is testing addresses (the **-bt** option) runs as root so that the routing is done in the same environment as a message delivery.

## 55.3 Running Exim without privilege

Some installations like to run Exim in an unprivileged state for more of its operation, for added security. Support for this mode of operation is provided by the global option **deliver\_drop\_privilege**. When this is set, the uid and gid are changed to the Exim user and group at the start of a delivery process (and also queue runner and address testing processes). This means that address routing is no longer run as root, and the deliveries themselves cannot change to any other uid.

Leaving the binary setuid to root, but setting **deliver\_drop\_privilege** means that the daemon can still be started in the usual way, and it can respond correctly to SIGHUP because the re-invocation regains root privilege.

An alternative approach is to make Exim setuid to the Exim user and also setgid to the Exim group. If you do this, the daemon must be started from a root process. (Calling Exim from a root process makes it behave in the way it does when it is setuid root.) However, the daemon cannot restart itself after a SIGHUP signal because it cannot regain privilege.

It is still useful to set **deliver\_drop\_privilege** in this case, because it stops Exim from trying to re-invoke itself to do a delivery after a message has been received. Such a re-invocation is a waste of resources because it has no effect.

If restarting the daemon is not an issue (for example, if **mua\_wrapper** is set, or *inetd* is being used instead of a daemon), having the binary setuid to the Exim user seems a clean approach, but there is one complication:

In this style of operation, Exim is running with the real uid and gid set to those of the calling process, and the effective uid/gid set to Exim's values. Ideally, any association with the calling process' uid/gid should be dropped, that is, the real uid/gid should be reset to the effective values so as to discard any privileges that the caller may have. While some operating systems have a function that permits this action for a non-root effective uid, quite a number of them do not. Because of this lack of standardization, Exim does not address this problem at this time.

For this reason, the recommended approach for “mostly unprivileged” running is to keep the Exim binary setuid to root, and to set **deliver\_drop\_privilege**. This also has the advantage of allowing a daemon to be used in the most straightforward way.

If you configure Exim not to run delivery processes as root, there are a number of restrictions on what you can do:

- You can deliver only as the Exim user/group. You should explicitly use the **user** and **group** options to override routers or local transports that normally deliver as the recipient. This makes sure that configurations that work in this mode function the same way in normal mode. Any implicit or explicit specification of another user causes an error.
- Use of *.forward* files is severely restricted, such that it is usually not worthwhile to include them in the configuration.
- Users who wish to use *.forward* would have to make their home directory and the file itself accessible to the Exim user. Pipe and append-to-file entries, and their equivalents in Exim filters, cannot be used. While they could be enabled in the Exim user's name, that would be insecure and not very useful.
- Unless the local user mailboxes are all owned by the Exim user (possible in some POP3 or IMAP-only environments):
  - (1) They must be owned by the Exim group and be writeable by that group. This implies you must set **mode** in the appendfile configuration, as well as the mode of the mailbox files themselves.
  - (2) You must set **no\_check\_owner**, since most or all of the files will not be owned by the Exim user.
  - (3) You must set **file\_must\_exist**, because Exim cannot set the owner correctly on a newly created mailbox when unprivileged. This also implies that new mailboxes need to be created manually.

These restrictions severely restrict what can be done in local deliveries. However, there are no restrictions on remote deliveries. If you are running a gateway host that does no local deliveries, setting **deliver\_drop\_privilege** gives more security at essentially no cost.

If you are using the **mua\_wrapper** facility (see chapter 51), **deliver\_drop\_privilege** is forced to be true.

## 55.4 Delivering to local files

Full details of the checks applied by *appendfile* before it writes to a file are given in chapter 26.

## 55.5 Running local commands

There are a number of ways in which an administrator can configure Exim to run commands based upon received, untrustworthy, data. Further, in some configurations a user who can control a *.forward* file can also arrange to run commands. Configuration to check includes, but is not limited to:

- Use of **use\_shell** in the pipe transport: various forms of shell command injection may be possible with this option present. It is dangerous and should be used only with considerable caution. Consider constraints which whitelist allowed characters in a variable which is to be used in a pipe transport that has **use\_shell** enabled.
- A number of options such as **forbid\_filter\_run**, **forbid\_filter\_perl**, **forbid\_filter\_dlfunc** and so forth which restrict facilities available to *.forward* files in a redirect router. If Exim is running on a central mail hub to which ordinary users do not have shell access, but home directories are NFS mounted (for instance) then administrators should review the list of these forbid options available, and should bear in mind that the options that may need forbidding can change as new features are added between releases.
- The **\${run...}** expansion item does not use a shell by default, but administrators can configure use of */bin/sh* as part of the command. Such invocations should be viewed with prejudicial suspicion.
- Administrators who use embedded Perl are advised to explore how Perl's taint checking might apply to their usage.
- Use of **\${expand...}** is somewhat analogous to shell's *eval* builtin and administrators are well advised to view its use with suspicion, in case (for instance) it allows a local-part to contain embedded Exim directives.
- Use of **\${match\_local\_part...}** and friends becomes more dangerous if Exim was built with **EXPAND\_LISTMATCH\_RHS** defined: the second string in each can reference arbitrary lists and files, rather than just being a list of opaque strings. The **EXPAND\_LISTMATCH\_RHS** option was added and set false by default because of real-world security vulnerabilities caused by its use with untrustworthy data injected in, for SQL injection attacks. Consider the use of the **inlisti** expansion condition instead.

## 55.6 Trust in configuration data

If configuration data for Exim can come from untrustworthy sources, there are some issues to be aware of:

- Use of **\${expand...}** may provide a path for shell injection attacks.
- Letting untrusted data provide a regular expression is unwise.
- Using **\${match...}** to apply a fixed regular expression against untrusted data may result in pathological behaviour within PCRE. Be aware of what "backtracking" means and consider options for being more strict with a regular expression. Avenues to explore include limiting what can match (avoiding *.* when *[a-z0-9]* or other character class will do), use of atomic grouping and possessive quantifiers or just not using regular expressions against untrusted data.
- It can be important to correctly use **\${quote:...}**, **\${quote\_local\_part:...}** and **\${quote\_<lookup-type>:...}** expansion items to ensure that data is correctly constructed.

- Some lookups might return multiple results, even though normal usage is only expected to yield one result.

## 55.7 IPv4 source routing

Many operating systems suppress IP source-routed packets in the kernel, but some cannot be made to do this, so Exim does its own check. It logs incoming IPv4 source-routed TCP calls, and then drops them. Things are all different in IPv6. No special checking is currently done.

## 55.8 The VRFY, EXPN, and ETRN commands in SMTP

Support for these SMTP commands is disabled by default. If required, they can be enabled by defining suitable ACLs.

## 55.9 Privileged users

Exim recognizes two sets of users with special privileges. Trusted users are able to submit new messages to Exim locally, but supply their own sender addresses and information about a sending host. For other users submitting local messages, Exim sets up the sender address from the uid, and doesn't permit a remote host to be specified.

However, an untrusted user is permitted to use the **-f** command line option in the special form **-f <>** to indicate that a delivery failure for the message should not cause an error report. This affects the message's envelope, but it does not affect the *Sender:* header. Untrusted users may also be permitted to use specific forms of address with the **-f** option by setting the **untrusted\_set\_sender** option.

Trusted users are used to run processes that receive mail messages from some other mail domain and pass them on to Exim for delivery either locally, or over the Internet. Exim trusts a caller that is running as root, as the Exim user, as any user listed in the **trusted\_users** configuration option, or under any group listed in the **trusted\_groups** option.

Admin users are permitted to do things to the messages on Exim's queue. They can freeze or thaw messages, cause them to be returned to their senders, remove them entirely, or modify them in various ways. In addition, admin users can run the Exim monitor and see all the information it is capable of providing, which includes the contents of files on the spool.

By default, the use of the **-M** and **-q** options to cause Exim to attempt delivery of messages on its queue is restricted to admin users. This restriction can be relaxed by setting the **no\_prod\_requires\_admin** option. Similarly, the use of **-bp** (and its variants) to list the contents of the queue is also restricted to admin users. This restriction can be relaxed by setting **no\_queue\_list\_requires\_admin**.

Exim recognizes an admin user if the calling process is running as root or as the Exim user or if any of the groups associated with the calling process is the Exim group. It is not necessary actually to be running under the Exim group. However, if admin users who are not root or the Exim user are to access the contents of files on the spool via the Exim monitor (which runs unprivileged), Exim must be built to allow group read access to its spool files.

By default, regular users are trusted to perform basic testing and introspection commands, as themselves. This setting can be tightened by setting the **commandline\_checks\_require\_admin** option. This affects most of the checking options, such as **-be** and anything else **-b\***.

## 55.10 Spool files

Exim's spool directory and everything it contains is owned by the Exim user and set to the Exim group. The mode for spool files is defined in the *Local/Makefile* configuration file, and defaults to 0640. This means that any user who is a member of the Exim group can access these files.

## 55.11 Use of argv[0]

Exim examines the last component of **argv[0]**, and if it matches one of a set of specific strings, Exim assumes certain options. For example, calling Exim with the last component of **argv[0]** set to “rsmtmp” is exactly equivalent to calling it with the option **-bS**. There are no security implications in this.

## 55.12 Use of %f formatting

The only use made of “%f” by Exim is in formatting load average values. These are actually stored in integer variables as 1000 times the load average. Consequently, their range is limited and so therefore is the length of the converted output.

## 55.13 Embedded Exim path

Exim uses its own path name, which is embedded in the code, only when it needs to re-exec in order to regain root privilege. Therefore, it is not root when it does so. If some bug allowed the path to get overwritten, it would lead to an arbitrary program’s being run as exim, not as root.

## 55.14 Dynamic module directory

Any dynamically loadable modules must be installed into the directory defined in `LOOKUP_MODULE_DIR` in *Local/Makefile* for Exim to permit loading it.

## 55.15 Use of sprintf()

A large number of occurrences of “sprintf” in the code are actually calls to *string\_sprintf()*, a function that returns the result in malloc’d store. The intermediate formatting is done into a large fixed buffer by a function that runs through the format string itself, and checks the length of each conversion before performing it, thus preventing buffer overruns.

The remaining uses of *sprintf()* happen in controlled circumstances where the output buffer is known to be sufficiently long to contain the converted string.

## 55.16 Use of debug\_printf() and log\_write()

Arbitrary strings are passed to both these functions, but they do their formatting by calling the function *string\_vformat()*, which runs through the format string itself, and checks the length of each conversion.

## 55.17 Use of strcat() and strcpy()

These are used only in cases where the output buffer is known to be large enough to hold the result.



## 56. Format of spool files

A message on Exim's queue consists of two files, whose names are the message id followed by -D and -H, respectively. The data portion of the message is kept in the -D file on its own. The message's envelope, status, and headers are all kept in the -H file, whose format is described in this chapter. Each of these two files contains the final component of its own name as its first line. This is insurance against disk crashes where the directory is lost but the files themselves are recoverable.

Some people are tempted into editing -D files in order to modify messages. You need to be extremely careful if you do this; it is not recommended and you are on your own if you do it. Here are some of the pitfalls:

- You must ensure that Exim does not try to deliver the message while you are fiddling with it. The safest way is to take out a write lock on the -D file, which is what Exim itself does, using *fcntl()*. If you update the file in place, the lock will be retained. If you write a new file and rename it, the lock will be lost at the instant of rename.
- If you change the number of lines in the file, the value of *\$body\_linecount*, which is stored in the -H file, will be incorrect and can cause incomplete transmission of messages or undeliverable messages.
- If the message is in MIME format, you must take care not to break it.
- If the message is cryptographically signed, any change will invalidate the signature.

All in all, modifying -D files is fraught with danger.

Files whose names end with -J may also be seen in the *input* directory (or its subdirectories when **split\_spool\_directory** is set). These are journal files, used to record addresses to which the message has been delivered during the course of a delivery attempt. If there are still undelivered recipients at the end, the -H file is updated, and the -J file is deleted. If, however, there is some kind of crash (for example, a power outage) before this happens, the -J file remains in existence. When Exim next processes the message, it notices the -J file and uses it to update the -H file before starting the next delivery attempt.

Files whose names end with -K or .eml may also be seen in the spool. These are temporaries used for DKIM or malware processing, when that is used. They should be tidied up by normal operations; any old ones are probably relics of crashes and can be removed.

### 56.1 Format of the -H file

The second line of the -H file contains the login name for the uid of the process that called Exim to read the message, followed by the numerical uid and gid. For a locally generated message, this is normally the user who sent the message. For a message received over TCP/IP via the daemon, it is normally the Exim user.

The third line of the file contains the address of the message's sender as transmitted in the envelope, contained in angle brackets. The sender address is empty for bounce messages. For incoming SMTP mail, the sender address is given in the MAIL command. For locally generated mail, the sender address is created by Exim from the login name of the current user and the configured **qualify\_domain**. However, this can be overridden by the **-f** option or a leading "From " line if the caller is trusted, or if the supplied address is "<>" or an address that matches **untrusted\_set\_senders**.

The fourth line contains two numbers. The first is the time that the message was received, in the conventional Unix form – the number of seconds since the start of the epoch. The second number is a count of the number of messages warning of delayed delivery that have been sent to the sender.

There follow a number of lines starting with a hyphen. These can appear in any order, and are omitted when not relevant:

**-acl** <number> <length>

This item is obsolete, and is not generated from Exim release 4.61 onwards; **-aclc** and **-aclm** are used instead. However, **-acl** is still recognized, to provide backward compatibility. In the old

format, a line of this form is present for every ACL variable that is not empty. The number identifies the variable; the **acl\_cx** variables are numbered 0–9 and the **acl\_mx** variables are numbered 10–19. The length is the length of the data string for the variable. The string itself starts at the beginning of the next line, and is followed by a newline character. It may contain internal newlines.

**-aclc** <rest-of-name> <length>

A line of this form is present for every ACL connection variable that is defined. Note that there is a space between **-aclc** and the rest of the name. The length is the length of the data string for the variable. The string itself starts at the beginning of the next line, and is followed by a newline character. It may contain internal newlines.

**-aclm** <rest-of-name> <length>

A line of this form is present for every ACL message variable that is defined. Note that there is a space between **-aclm** and the rest of the name. The length is the length of the data string for the variable. The string itself starts at the beginning of the next line, and is followed by a newline character. It may contain internal newlines.

**-active\_hostname** <hostname>

This is present if, when the message was received over SMTP, the value of *\$smtp\_active\_hostname* was different to the value of *\$primary\_hostname*.

**-allow\_unqualified\_recipient**

This is present if unqualified recipient addresses are permitted in header lines (to stop such addresses from being qualified if rewriting occurs at transport time). Local messages that were input using **-bnq** and remote messages from hosts that match **recipient\_unqualified\_hosts** set this flag.

**-allow\_unqualified\_sender**

This is present if unqualified sender addresses are permitted in header lines (to stop such addresses from being qualified if rewriting occurs at transport time). Local messages that were input using **-bnq** and remote messages from hosts that match **sender\_unqualified\_hosts** set this flag.

**-auth\_id** <text>

The id information for a message received on an authenticated SMTP connection – the value of the *\$authenticated\_id* variable.

**-auth\_sender** <address>

The address of an authenticated sender – the value of the *\$authenticated\_sender* variable.

**-body\_linecount** <number>

This records the number of lines in the body of the message, and is present unless **-spool\_file\_wireformat** is.

**-body\_zerocount** <number>

This records the number of binary zero bytes in the body of the message, and is present if the number is greater than zero.

**-deliver\_firsttime**

This is written when a new message is first added to the spool. When the spool file is updated after a deferral, it is omitted.

**-frozen** <time>

The message is frozen, and the freezing happened at <time>.

**-helo\_name** <text>

This records the host name as specified by a remote host in a HELO or EHLO command.

**-host\_address** <address>.<port>

This records the IP address of the host from which the message was received and the remote port number that was used. It is omitted for locally generated messages.

**-host\_auth** <text>

If the message was received on an authenticated SMTP connection, this records the name of the authenticator – the value of the *\$sender\_host\_authenticated* variable.

**-host\_lookup\_failed**

This is present if an attempt to look up the sending host's name from its IP address failed. It corresponds to the *\$host\_lookup\_failed* variable.

**-host\_name** <text>

This records the name of the remote host from which the message was received, if the host name was looked up from the IP address when the message was being received. It is not present if no reverse lookup was done.

**-ident** <text>

For locally submitted messages, this records the login of the originating user, unless it was a trusted user and the **-oMt** option was used to specify an ident value. For messages received over TCP/IP, this records the ident string supplied by the remote host, if any.

**-interface\_address** <address>.<port>

This records the IP address of the local interface and the port number through which a message was received from a remote host. It is omitted for locally generated messages.

**-local**

The message is from a local sender.

**-localerror**

The message is a locally-generated bounce message.

**-local\_scan** <string>

This records the data string that was returned by the *local\_scan()* function when the message was received – the value of the *\$local\_scan\_data* variable. It is omitted if no data was returned.

**-manual\_thaw**

The message was frozen but has been thawed manually, that is, by an explicit Exim command rather than via the auto-thaw process.

**-N**

A testing delivery process was started using the **-N** option to suppress any actual deliveries, but delivery was deferred. At any further delivery attempts, **-N** is assumed.

**-received\_protocol**

This records the value of the *\$received\_protocol* variable, which contains the name of the protocol by which the message was received.

**-sender\_set\_untrusted**

The envelope sender of this message was set by an untrusted local caller (used to ensure that the caller is displayed in queue listings).

**-spam\_score\_int** <number>

If a message was scanned by SpamAssassin, this is present. It records the value of *\$spam\_score\_int*.

**-spool\_file\_wireformat**

The **-D** file for this message is in wire-format (for ESMTP CHUNKING) rather than Unix-format. The line-ending is CRLF rather than newline. There is still, however, no leading-dot-stuffing.

**-tls\_certificate\_verified**

A TLS certificate was received from the client that sent this message, and the certificate was verified by the server.

**-tls\_cipher** <cipher name>

When the message was received over an encrypted connection, this records the name of the cipher suite that was used.

**-tls\_peerdn** <peer DN>

When the message was received over an encrypted connection, and a certificate was received from the client, this records the Distinguished Name from that certificate.

Following the options there is a list of those addresses to which the message is not to be delivered. This set of addresses is initialized from the command line when the **-t** option is used and **extract\_**

**addresses\_remove\_arguments** is set; otherwise it starts out empty. Whenever a successful delivery is made, the address is added to this set. The addresses are kept internally as a balanced binary tree, and it is a representation of that tree which is written to the spool file. If an address is expanded via an alias or forward file, the original address is added to the tree when deliveries to all its child addresses are complete.

If the tree is empty, there is a single line in the spool file containing just the text “XX”. Otherwise, each line consists of two letters, which are either Y or N, followed by an address. The address is the value for the node of the tree, and the letters indicate whether the node has a left branch and/or a right branch attached to it, respectively. If branches exist, they immediately follow. Here is an example of a three-node tree:

```
YY darcy@austen.fict.example
NN alice@wonderland.fict.example
NN editor@thesaurus.ref.example
```

After the non-recipients tree, there is a list of the message’s recipients. This is a simple list, preceded by a count. It includes all the original recipients of the message, including those to whom the message has already been delivered. In the simplest case, the list contains one address per line. For example:

```
4
editor@thesaurus.ref.example
darcy@austen.fict.example
rdo@foundation
alice@wonderland.fict.example
```

However, when a child address has been added to the top-level addresses as a result of the use of the **one\_time** option on a *redirect* router, each line is of the following form:

*<top-level address> <errors\_to address> <length>,<parent number>#<flag bits>*

The 01 flag bit indicates the presence of the three other fields that follow the top-level address. Other bits may be used in future to support additional fields. The *<parent number>* is the offset in the recipients list of the original parent of the “one time” address. The first two fields are the envelope sender that is associated with this address and its length. If the length is zero, there is no special envelope sender (there are then two space characters in the line). A non-empty field can arise from a *redirect* router that has an **errors\_to** setting.

A blank line separates the envelope and status information from the headers which follow. A header may occupy several lines of the file, and to save effort when reading it in, each header is preceded by a number and an identifying character. The number is the number of characters in the header, including any embedded newlines and the terminating newline. The character is one of the following:

<i>&lt;blank&gt;</i>	header in which Exim has no special interest
B	<i>Bcc:</i> header
C	<i>Cc:</i> header
F	<i>From:</i> header
I	<i>Message-id:</i> header
P	<i>Received:</i> header – P for “postmark”
R	<i>Reply-To:</i> header
S	<i>Sender:</i> header
T	<i>To:</i> header
*	replaced or deleted header

Deleted or replaced (rewritten) headers remain in the spool file for debugging purposes. They are not transmitted when the message is delivered. Here is a typical set of headers:

```
111P Received: by hobbit.fict.example with local (Exim 4.00)
id 14y9EI-00026G-00; Fri, 11 May 2001 10:28:59 +0100
049 Message-Id: <E14y9EI-00026G-00@hobbit.fict.example>
038* X-rewrote-sender: bb@hobbit.fict.example
042* From: Bilbo Baggins <bb@hobbit.fict.example>
049F From: Bilbo Baggins <B.Baggins@hobbit.fict.example>
```

```
099* To: alice@wonderland.fict.example, rdo@foundation,  
darcy@austen.fict.example, editor@thesaurus.ref.example  
104T To: alice@wonderland.fict.example, rdo@foundation.example,  
darcy@austen.fict.example, editor@thesaurus.ref.example  
038 Date: Fri, 11 May 2001 10:28:59 +0100
```

The asterisked headers indicate that the envelope sender, *From:* header, and *To:* header have been rewritten, the last one because routing expanded the unqualified domain *foundation*.

## 56.2 Format of the -D file

The data file is traditionally in Unix-standard format: lines are ended with an ASCII newline character. However, when the **spool\_wireformat** main option is used some -D files can have an alternate format. This is flagged by a **-spool\_file\_wireformat** line in the corresponding -H file. The -D file lines (not including the first name-component line) are suitable for direct copying to the wire when transmitting using the ESMTP CHUNKING option, meaning lower processing overhead. Lines are terminated with an ASCII CRLF pair. There is no dot-stuffing (and no dot-termination).

## 57. Support for DKIM (DomainKeys Identified Mail)

DKIM is a mechanism by which messages sent by some entity can be provably linked to a domain which that entity controls. It permits reputation to be tracked on a per-domain basis, rather than merely upon source IP address. DKIM is documented in RFC 4871.

DKIM support is compiled into Exim by default if TLS support is present. It can be disabled by setting `DISABLE_DKIM=yes` in *Local/Makefile*.

Exim's DKIM implementation allows for

- (1) Signing outgoing messages: This function is implemented in the SMTP transport. It can co-exist with all other Exim features (including transport filters) except cutthrough delivery.
- (2) Verifying signatures in incoming messages: This is implemented by an additional ACL (`acl_smtp_dkim`), which can be called several times per message, with different signature contexts.

In typical Exim style, the verification implementation does not include any default "policy". Instead it enables you to build your own policy using Exim's standard controls.

Please note that verification of DKIM signatures in incoming mail is turned on by default for logging purposes. For each signature in incoming email, exim will log a line displaying the most important signature details, and the signature status. Here is an example (with line-breaks added for clarity):

```
2009-09-09 10:22:28 1M1IRf-0003LU-U3 DKIM:
d=facebookmail.com s=q1-2009b
c=relaxed/relaxed a=rsa-sha1
i=@facebookmail.com t=1252484542 [verification succeeded]
```

You might want to turn off DKIM verification processing entirely for internal or relay mail sources. To do that, set the **`dkim_disable_verify`** ACL control modifier. This should typically be done in the RCPT ACL, at points where you accept mail from relay sources (internal hosts or authenticated senders).

### 57.1 Signing outgoing messages

Signing is enabled by setting private options on the SMTP transport. These options take (expandable) strings as arguments.

<b><code>dkim_domain</code></b>	Use: <i>smtp</i>	Type: <i>string</i>	Default: <i>list†</i>
---------------------------------	------------------	---------------------	-----------------------

The domain(s) you want to sign with.

After expansion, this can be a list. Each element in turn is put into the **`$dkim_domain`** expansion variable while expanding the remaining signing options.

If it is empty after expansion, DKIM signing is not done.

<b><code>dkim_selector</code></b>	Use: <i>smtp</i>	Type: <i>string</i>	Default: <i>list†</i>
-----------------------------------	------------------	---------------------	-----------------------

This sets the key selector string.

After expansion, which can use *`$dkim_domain`*, this can be a list. Each element in turn is put in the expansion variable **`$dkim_selector`** which may be used in the **`dkim_private_key`** option along with **`$dkim_domain`**. If the option is empty after expansion, DKIM signing is not done for this domain.

<b><code>dkim_private_key</code></b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------------------------	------------------	----------------------	-----------------------

This sets the private key to use. You can use the **`$dkim_domain`** and **`$dkim_selector`** expansion variables to determine the private key to use. The result can either

- be a valid RSA private key in ASCII armor, including line breaks.
- start with a slash, in which case it is treated as a file that contains the private key.
- be "0", "false" or the empty string, in which case the message will not be signed. This case will not result in an error, even if **dkim\_strict** is set.

If the option is empty after expansion, DKIM signing is not done.

<b>dkim_hash</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>sha256</i>
------------------	------------------	----------------------	------------------------

Can be set alternatively to "sha1" to use an alternate hash method. Note that sha1 is now considered insecure, and deprecated.

<b>dkim_identity</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
----------------------	------------------	----------------------	-----------------------

If set after expansion, the value is used to set an "i=" tag in the signing header. The DKIM standards restrict the permissible syntax of this optional tag to a mail address, with possibly-empty local part, an @, and a domain identical to or subdomain of the "d=" tag value. Note that Exim does not check the value.

<b>dkim_canon</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
-------------------	------------------	----------------------	-----------------------

This option sets the canonicalization method used when signing a message. The DKIM RFC currently supports two methods: "simple" and "relaxed". The option defaults to "relaxed" when unset. Note: the current implementation only supports using the same canonicalization method for both headers and body.

<b>dkim_strict</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>unset</i>
--------------------	------------------	----------------------	-----------------------

This option defines how Exim behaves when signing a message that should be signed fails for some reason. When the expansion evaluates to either "1" or "true", Exim will defer. Otherwise Exim will send the message unsigned. You can use the **\$dkim\_domain** and **\$dkim\_selector** expansion variables here.

<b>dkim_sign_headers</b>	Use: <i>smtp</i>	Type: <i>string†</i>	Default: <i>see below</i>
--------------------------	------------------	----------------------	---------------------------

If set, this option must expand to a colon-separated list of header names.

Headers with these names, or the absence of such a header, will be included in the message signature. When unspecified, the header names listed in RFC4871 will be used, whether or not each header is present in the message. The default list is available for the expansion in the macro "\_DKIM\_SIGN\_HEADERS".

If a name is repeated, multiple headers by that name (or the absence thereof) will be signed. The textually later headers in the headers part of the message are signed first, if there are multiples.

A name can be prefixed with either an '=' or a '+' character. If an '=' prefix is used, all headers that are present with this name will be signed. If a '+' prefix is used, all headers that are present with this name will be signed, and one signature added for a missing header with the name will be appended.

## 57.2 Verifying DKIM signatures in incoming mail

Verification of DKIM signatures in SMTP incoming email is implemented via the **acl\_smtp\_dkim** ACL. By default, this ACL is called once for each syntactically(!) correct signature in the incoming message. A missing ACL definition defaults to accept. If any ACL call does not accept, the message is not accepted. If a cutthrough delivery was in progress for the message, that is summarily dropped (having wasted the transmission effort).

To evaluate the signature in the ACL a large number of expansion variables containing the signature status and its details are set up during the runtime of the ACL.

Calling the ACL only for existing signatures is not sufficient to build more advanced policies. For that reason, the global option **dkim\_verify\_signers**, and a global expansion variable **\$dkim\_signers** exist.

The global option **dkim\_verify\_signers** can be set to a colon-separated list of DKIM domains or identities for which the ACL **acl\_smtp\_dkim** is called. It is expanded when the message has been received. At this point, the expansion variable **\$dkim\_signers** already contains a colon-separated list of signer domains and identities for the message. When **dkim\_verify\_signers** is not specified in the main configuration, it defaults as:

```
dkim_verify_signers = $dkim_signers
```

This leads to the default behaviour of calling **acl\_smtp\_dkim** for each DKIM signature in the message. Current DKIM verifiers may want to explicitly call the ACL for known domains or identities. This would be achieved as follows:

```
dkim_verify_signers = paypal.com:ebay.com:$dkim_signers
```

This would result in **acl\_smtp\_dkim** always being called for "paypal.com" and "ebay.com", plus all domains and identities that have signatures in the message. You can also be more creative in constructing your policy. For example:

```
dkim_verify_signers = $sender_address_domain:$dkim_signers
```

If a domain or identity is listed several times in the (expanded) value of **dkim\_verify\_signers**, the ACL is only called once for that domain or identity.

If multiple signatures match a domain (or identity), the ACL is called once for each matching signature.

Inside the **acl\_smtp\_dkim**, the following expansion variables are available (from most to least important):

#### **\$dkim\_cur\_signer**

The signer that is being evaluated in this ACL run. This can be a domain or an identity. This is one of the list items from the expanded main option **dkim\_verify\_signers** (see above).

#### **\$dkim\_verify\_status**

Within the DKIM ACL, a string describing the general status of the signature. One of

- **none**: There is no signature in the message for the current domain or identity (as reflected by **\$dkim\_cur\_signer**).
- **invalid**: The signature could not be verified due to a processing error. More detail is available in **\$dkim\_verify\_reason**.
- **fail**: Verification of the signature failed. More detail is available in **\$dkim\_verify\_reason**.
- **pass**: The signature passed verification. It is valid.

This variable can be overwritten using an ACL 'set' modifier. This might, for instance, be done to enforce a policy restriction on hash-method or key-size:

```
warn␣condition =␣${if eq {$dkim_algo}{rsa-sha1}}  
␣condition =␣${if eq {$dkim_verify_status}{pass}}  
␣logwrite =␣NOTE: forcing dkim verify fail (was pass)  
␣set dkim_verify_status = fail  
␣set dkim_verify_reason = hash too weak
```

After all the DKIM ACL runs have completed, the value becomes a colon-separated list of the values after each run.

#### **\$dkim\_verify\_reason**

A string giving a little bit more detail when **\$dkim\_verify\_status** is either "fail" or "invalid". One of



- **pubkey\_unavailable** (when `$dkim_verify_status="invalid"`): The public key for the domain could not be retrieved. This may be a temporary problem.
- **pubkey\_syntax** (when `$dkim_verify_status="invalid"`): The public key record for the domain is syntactically invalid.
- **bodyhash\_mismatch** (when `$dkim_verify_status="fail"`): The calculated body hash does not match the one specified in the signature header. This means that the message body was modified in transit.
- **signature\_incorrect** (when `$dkim_verify_status="fail"`): The signature could not be verified. This may mean that headers were modified, re-written or otherwise changed in a way which is incompatible with DKIM verification. It may of course also mean that the signature is forged.

This variable can be overwritten using an ACL 'set' modifier.

### **\$dkim\_domain**

The signing domain. IMPORTANT: This variable is only populated if there is an actual signature in the message for the current domain or identity (as reflected by `$dkim_cur_signer`).

### **\$dkim\_identity**

The signing identity, if present. IMPORTANT: This variable is only populated if there is an actual signature in the message for the current domain or identity (as reflected by `$dkim_cur_signer`).

### **\$dkim\_selector**

The key record selector string.

### **\$dkim\_algo**

The algorithm used. One of 'rsa-sha1' or 'rsa-sha256'.

### **\$dkim\_canon\_body**

The body canonicalization method. One of 'relaxed' or 'simple'.

### **\$dkim\_canon\_headers**

The header canonicalization method. One of 'relaxed' or 'simple'.

### **\$dkim\_copiedheaders**

A transcript of headers and their values which are included in the signature (copied from the 'z=' tag of the signature). Note that RFC6376 requires that verification fail if the From: header is not included in the signature. Exim does not enforce this; sites wishing strict enforcement should code the check explicitly.

### **\$dkim\_bodylength**

The number of signed body bytes. If zero ("0"), the body is unsigned. If no limit was set by the signer, "9999999999999999" is returned. This makes sure that this variable always expands to an integer value.

### **\$dkim\_created**

UNIX timestamp reflecting the date and time when the signature was created. When this was not specified by the signer, "0" is returned.

### **\$dkim\_expires**

UNIX timestamp reflecting the date and time when the signer wants the signature to be treated as "expired". When this was not specified by the signer, "9999999999999999" is returned. This makes it possible to do useful integer size comparisons against this value.

### **\$dkim\_headernames**

A colon-separated list of names of headers included in the signature.

### **\$dkim\_key\_testing**

"1" if the key record has the "testing" flag set, "0" if not.

### **\$dkim\_key\_nosubdomains**

"1" if the key record forbids subdomaining, "0" otherwise.

### **\$dkim\_key\_srvtype**

Service type (tag s=) from the key record. Defaults to "\*" if not specified in the key record.

**\$dkim\_key\_granularity**

Key granularity (tag g=) from the key record. Defaults to "\*" if not specified in the key record.

**\$dkim\_key\_notes**

Notes from the key record (tag n=).

**\$dkim\_key\_length**

Number of bits in the key.

In addition, two ACL conditions are provided:

**dkim\_signers**

ACL condition that checks a colon-separated list of domains or identities for a match against the domain or identity that the ACL is currently verifying (reflected by **\$dkim\_cur\_signer**). This is typically used to restrict an ACL verb to a group of domains or identities. For example:

```
# Warn when Mail purportedly from GMail has no gmail signature
warn log_message = GMail sender without gmail.com DKIM signature
  sender_domains = gmail.com
  dkim_signers = gmail.com
  dkim_status = none
```

Note that the above does not check for a total lack of DKIM signing; for that check for empty *\$h\_DKIM-Signature:* in the data ACL.

**dkim\_status**

ACL condition that checks a colon-separated list of possible DKIM verification results against the actual result of verification. This is typically used to restrict an ACL verb to a list of verification outcomes, for example:

```
deny message = Mail from Paypal with invalid/missing signature
  sender_domains = paypal.com:paypal.de
  dkim_signers = paypal.com:paypal.de
  dkim_status = none:invalid:fail
```

The possible status keywords are: 'none', 'invalid', 'fail' and 'pass'. Please see the documentation of the **\$dkim\_verify\_status** expansion variable above for more information of what they mean.

## 58. Proxies

A proxy is an intermediate system through which communication is passed. Proxies may provide a security, availability or load-distribution function.

### 58.1 Inbound proxies

Exim has support for receiving inbound SMTP connections via a proxy that uses “Proxy Protocol” to speak to it. To include this support, include “`SUPPORT_PROXY=yes`” in Local/Makefile.

It was built on specifications from: (<http://haproxy.1wt.eu/download/1.5/doc/proxy-protocol.txt>). That URL was revised in May 2014 to version 2 spec: (<http://git.1wt.eu/web?p=haproxy.git;a=commitdiff;h=afb768340c9d7e50d8e>).

The purpose of this facility is so that an application load balancer, such as HAProxy, can sit in front of several Exim servers to distribute load. Exim uses the local protocol communication with the proxy to obtain the remote SMTP system IP address and port information. There is no logging if a host passes or fails Proxy Protocol negotiation, but it can easily be determined and recorded in an ACL (example is below).

Use of a proxy is enabled by setting the **hosts\_proxy** main configuration option to a hostlist; connections from these hosts will use Proxy Protocol. Exim supports both version 1 and version 2 of the Proxy Protocol and automatically determines which version is in use.

The Proxy Protocol header is the first data received on a TCP connection and is inserted before any TLS-on-connect handshake from the client; Exim negotiates TLS between Exim-as-server and the remote client, not between Exim and the proxy server.

The following expansion variables are usable (“internal” and “external” here refer to the interfaces of the proxy):

<i>proxy_external_address</i>	IP of host being proxied or IP of remote interface of proxy
<i>proxy_external_port</i>	Port of host being proxied or Port on remote interface of proxy
<i>proxy_local_address</i>	IP of proxy server inbound or IP of local interface of proxy
<i>proxy_local_port</i>	Port of proxy server inbound or Port on local interface of proxy
<i>proxy_session</i>	boolean: SMTP connection via proxy

If *\$proxy\_session* is set but *\$proxy\_external\_address* is empty there was a protocol error.

Since the real connections are all coming from the proxy, and the per host connection tracking is done before Proxy Protocol is evaluated, **smtp\_accept\_max\_per\_host** must be set high enough to handle all of the parallel volume you expect per inbound proxy. With the option set so high, you lose the ability to protect your server from many connections from one IP. In order to prevent your server from overload, you need to add a per connection ratelimit to your connect ACL. A possible solution is:

```
# Set max number of connections per host
LIMIT = 5
# Or do some kind of IP lookup in a flat file or database
# LIMIT = ${lookup{$sender_host_address}iplsearch{/etc/exim/proxy_limits}}

defer message      = Too many connections from this IP right now
  ratelimit        = LIMIT / 5s / per_conn / strict
```

### 58.2 Outbound proxies

Exim has support for sending outbound SMTP via a proxy using a protocol called SOCKS5 (defined by RFC1928). The support can be optionally included by defining `SUPPORT_SOCKS=yes` in Local/Makefile.

Use of a proxy is enabled by setting the **socks\_proxy** option on an smtp transport. The option value is expanded and should then be a list (colon-separated by default) of proxy specifiers. Each proxy

specifier is a list (space-separated by default) where the initial element is an IP address and any subsequent elements are options.

Options are a string <name>=<value>. The list of options is in the following table:

<i>auth</i>	authentication method
<i>name</i>	authentication username
<i>pass</i>	authentication password
<i>port</i>	tcp port
<i>tmo</i>	connection timeout
<i>pri</i>	priority
<i>weight</i>	selection bias

More details on each of these options follows:

- **auth**: Either “none” (default) or “name”. Using “name” selects username/password authentication per RFC 1929 for access to the proxy. Default is “none”.
- **name**: sets the username for the “name” authentication method. Default is empty.
- **pass**: sets the password for the “name” authentication method. Default is empty.
- **port**: the TCP port number to use for the connection to the proxy. Default is 1080.
- **tmo**: sets a connection timeout in seconds for this proxy. Default is 5.
- **pri**: specifies a priority for the proxy within the list, higher values being tried first. The default priority is 1.
- **weight**: specifies a selection bias. Within a priority set servers are queried in a random fashion, weighted by this value. The default value for selection bias is 1.

Proxies from the list are tried according to their priority and weight settings until one responds. The timeout for the overall connection applies to the set of proxied attempts.

## 58.3 Logging

To log the (local) IP of a proxy in the incoming or delivery logline, add “+proxy” to the **log\_selector** option. This will add a component tagged with “PRX=” to the line.

## 59. Internationalisation

Exim has support for Internationalised mail names. To include this it must be built with `SUPPORT_I18N` and the `libidn` library. Standards supported are RFCs 2060, 5890, 6530 and 6533.

If Exim is built with `SUPPORT_I18N_2008` (in addition to `SUPPORT_I18N`, not instead of it) then IDNA2008 is supported; this adds an extra library requirement, upon `libidn2`.

### 59.1 MTA operations

The main configuration option `smtputf8_advertise_hosts` specifies a host list. If this matches the sending host and `accept_8bitmime` is true (the default) then the ESMTP option `SMTPUTF8` will be advertised.

If the sender specifies the `SMTPUTF8` option on a `MAIL` command international handling for the message is enabled and the expansion variable `$message_smtputf8` will have value `TRUE`.

The option `allow_utf8_domains` is set to true for this message. All DNS lookups are converted to a-label form whatever the setting of `allow_utf8_domains` when Exim is built with `SUPPORT_I18N`.

Both localparts and domain are maintained as the original UTF-8 form internally; any comparison or regular-expression use will require appropriate care. Filenames created, eg. by the appendfile transport, will have UTF-8 names.

HELO names sent by the smtp transport will have any UTF-8 components expanded to a-label form, and any certificate name checks will be done using the a-label form of the name.

Log lines and Received-by: header lines will acquire a "utf8" prefix on the protocol element, eg. `utf8smtp`.

The following expansion operators can be used:

```
{utf8_domain_to_alabel:str}
{utf8_domain_from_alabel:str}
{utf8_localpart_to_alabel:str}
{utf8_localpart_from_alabel:str}
```

ACLs may use the following modifier:

```
control = utf8_downconvert
control = utf8_downconvert/<value>
```

This sets a flag requiring that addresses are converted to a-label form before smtp delivery, for use in a Message Submission Agent context. If a value is appended it may be:

- 1 (default) mandatory downconversion
- 0 no downconversion
- 1 if `SMTPUTF8` not supported by destination host

If `mua_wrapper` is set, the `utf8_downconvert` control is initially set to -1.

There is no explicit support for `VRIFY` and `EXPN`. Configurations supporting these should inspect `$smtp_command_argument` for an `SMTPUTF8` argument.

There is no support for LMTP on Unix sockets. Using the "lmtp" protocol option on an smtp transport, for LMTP over TCP, should work as expected.

There is no support for DSN unitext handling, and no provision for converting logging from or to UTF-8.

### 59.2 MDA operations

To aid in constructing names suitable for IMAP folders the following expansion operator can be used:

```
{imapfolder {<string>} {<sep>} {<specials>}}
```

The string is converted from the charset specified by the "headers charset" command (in a filter file) or **headers\_charset** main configuration option (otherwise), to the modified UTF-7 encoding specified by RFC 2060, with the following exception: All occurrences of <sep> (which has to be a single character) are replaced with periods ("."), and all periods and slashes that are not <sep> and are not in the <specials> string are BASE64 encoded.

The third argument can be omitted, defaulting to an empty string. The second argument can be omitted, defaulting to "/".

This is the encoding used by Courier for Maildir names on disk, and followed by many other IMAP servers.

Examples:

```
{imapfolder {Foo/Bar}}          yields Foo.Bar
${imapfolder {Foo/Bar}}{.}/{/}} yields Foo&AC8-Bar
${imapfolder {Räksmörgås}}      yields R&AOQ-ksm&APY-rg&AOU-s
```

Note that the source charset setting is vital, and also that characters must be representable in UTF-16.

## 60. Events

The events mechanism in Exim can be used to intercept processing at a number of points. It was originally invented to give a way to do customised logging actions (for example, to a database) but can also be used to modify some processing actions.

Most installations will never need to use Events. The support can be left out of a build by defining `DISABLE_EVENT=yes` in *Local/Makefile*.

There are two major classes of events: main and transport. The main configuration option **event\_action** controls reception events; a transport option **event\_action** controls delivery events.

Both options are a string which is expanded when the event fires. An example might look like:

```
event_action = ${if eq {msg:delivery}{$event_name} \
${lookup pgsql {SELECT * FROM record_Delivery( \
  '${quote_pgsql:$sender_address_domain}', \
  '${quote_pgsql:${lc:$sender_address_local_part}}', \
  '${quote_pgsql:$domain}', \
  '${quote_pgsql:${lc:$local_part}}', \
  '${quote_pgsql:$host_address}', \
  '${quote_pgsql:${lc:$host}}', \
  '${quote_pgsql:$message_exim_id}')} \
} {} }
```

Events have names which correspond to the point in process at which they fire. The name is placed in the variable *\$event\_name* and the event action expansion must check this, as it will be called for every possible event type.

The current list of events is:

msg:complete	after	main	per message
msg:delivery	after	transport	per recipient
msg:rcpt:host:defer	after	transport	per recipient per host
msg:rcpt:defer	after	transport	per recipient
msg:host:defer	after	transport	per attempt
msg:fail:delivery	after	transport	per recipient
msg:fail:internal	after	main	per recipient
tcp:connect	before	transport	per connection
tcp:close	after	transport	per connection
tls:cert	before	both	per certificate in verification chain
smtp:connect	after	transport	per connection

New event types may be added in future.

The event name is a colon-separated list, defining the type of event in a tree of possibilities. It may be used as a list or just matched on as a whole. There will be no spaces in the name.

The second column in the table above describes whether the event fires before or after the action is associates with. Those which fire before can be used to affect that action (more on this below).

The third column in the table above says what section of the configuration should define the event action.

An additional variable, *\$event\_data*, is filled with information varying with the event type:

msg:delivery	smtp confirmation message
msg:rcpt:host:defer	error string
msg:rcpt:defer	error string
msg:host:defer	error string
tls:cert	verification chain depth
smtp:connect	smtp banner

The :defer events populate one extra variable: *\$event\_defer\_errno*.

For complex operations an ACL expansion can be used in **event\_action** however due to the multiple contexts that Exim operates in during the course of its processing:

- variables set in transport events will not be visible outside that transport call
- `acl_m` variables in a server context are lost on a new connection, and after `smtp helo/ehlo/mail/starttls/rset` commands

Using an ACL expansion with the `logwrite` modifier can be a useful way of writing to the main log.

The expansion of the `event_action` option should normally return an empty string. Should it return anything else the following will be forced:

<code>msg:delivery</code>	(ignored)
<code>msg:host:defer</code>	(ignored)
<code>msg:fail:delivery</code>	(ignored)
<code>tcp:connect</code>	do not connect
<code>tcp:close</code>	(ignored)
<code>tls:cert</code>	refuse verification
<code>smtp:connect</code>	close connection

No other use is made of the result string.

For a `tcp:connect` event, if the connection is being made to a proxy then the address and port variables will be that of the proxy and not the target system.

For `tls:cert` events, if GnuTLS is in use this will trigger only per chain element received on the connection. For OpenSSL it will trigger for every chain element including those loaded locally.



## 61. Adding new drivers or lookup types

The following actions have to be taken in order to add a new router, transport, authenticator, or lookup type to Exim:

- (1) Choose a name for the driver or lookup type that does not conflict with any existing name; I will use “newdriver” in what follows.
- (2) Add to *src/EDITME* the line:

```
<type>_NEWDRIVER=yes
```

where *<type>* is ROUTER, TRANSPORT, AUTH, or LOOKUP. If the code is not to be included in the binary by default, comment this line out. You should also add any relevant comments about the driver or lookup type.

- (3) Add to *src/config.h.defaults* the line:

```
#define <type>_NEWDRIVER
```

- (4) Edit *src/drtables.c*, adding conditional code to pull in the private header and create a table entry as is done for all the other drivers and lookup types.
- (5) Edit *scripts/lookups-Makefile* if this is a new lookup; there is a for-loop near the bottom, ranging the *name\_mod* variable over a list of all lookups. Add your NEWDRIVER to that list. As long as the dynamic module would be named *newdriver.so*, you can use the simple form that most lookups have.
- (6) Edit *Makefile* in the appropriate sub-directory (*src/routers*, *src/transport*s, *src/auth*s, or *src/lookups*); add a line for the new driver or lookup type and add it to the definition of OBJ.
- (7) Create *newdriver.h* and *newdriver.c* in the appropriate sub-directory of *src*.
- (8) Edit *scripts/MakeLinks* and add commands to link the *.h* and *.c* files as for other drivers and lookups.

Then all you need to do is write the code! A good way to start is to make a proforma by copying an existing module of the same type, globally changing all occurrences of the name, and cutting out most of the code. Note that any options you create must be listed in alphabetical order, because the tables are searched using a binary chop procedure.

There is a *README* file in each of the sub-directories of *src* describing the interface that is expected.

# Options index

## Symbols

-- 29  
--help 29  
--version 29  
-Ac 29  
-Am 29  
-B 29  
-bd 29  
-bdf 30  
-be 30, 102, 470  
-bem 30, 102  
-bF 30, 470  
-bf 30, 470  
-bfd 31  
-bfl 31  
-bfp 31  
-bfs 31  
-bh 31, 359, 470  
-bhc 32  
-bi 32, 163  
-bI:dscp 32  
-bI:help 32  
-bI:sieve 32  
-bm 32  
-bmalware 33  
-bnq 33  
-bP 33  
-bp 34, 186  
-bpa 34  
-bpc 34  
-bpr 34  
-bpra 34  
-bpru 35  
-bpu 35  
-brt 35  
-brw 35  
-bS 35  
-bs 35  
-bt 36, 207  
-bV 36  
-bv 36, 220, 470  
-bvs 37  
-bw 37  
-C 37, 470  
-D 38, 470  
-d 38  
-dd 39  
-dropcr 39  
-E 39  
-ex 40  
-F 40  
-f 40, 473  
    for address testing 36  
    for filter testing 31  
    overriding "From" line 33  
-G 40  
-h 40  
-i 40  
-L 41  
-M 41, 186, 473  
-m 43  
-Mar 41  
-MC 41  
-Mc 42  
-MCA 41  
-MCD 41  
-MCG 41  
-MCK 41  
-MCP 41  
-MCQ 41  
-MCS 42  
-MCT 42  
-MCt 42  
-Mes 42  
-Mf 42  
-Mg 42  
-Mmad 42  
-Mmd 42  
-Mrm 42  
-Mset 43  
-Mt 43  
-Mvb 43  
-Mvc 43  
-Mvh 43  
-Mvl 43  
-N 43  
-n 43  
-O 43  
-oA 43  
-oB 43  
-odb 43  
-odf 44  
-odi 44  
-odq 44  
-odqs 44  
-oee 44  
-oem 44  
-oep 45  
-oeq 45  
-oew 45  
-oi 45  
-oittrue 45  
-om 46  
-oMa 45  
-oMaa 45  
-oMai 45  
-oMas 45  
-oMi 45  
-oMm 46  
-oMr 46  
-oMs 46  
-oMt 46  
-oo 46  
-oP 46  
-or 46  
-os 46, 196  
-ov 46  
-oX 46  
-p 47  
-pd 47  
-ps 47  
-q 47, 186, 473  
-qf 48  
-qff 48

**-qG** 48  
**-qi** 48  
**-ql** 48  
**-qq** 47  
**-qR** 48  
**-qS** 48  
**-R** 49, 186  
**-r** 49  
**-S** 49  
**-t** 49, 172  
**-ti** 50  
**-tls-on-connect** 50  
**-Tqt** 49  
**-U** 50  
**-v** 50  
**-X** 50  
**-x** 50  
**-z** 50

## A

**accept\_8bitmime** 159  
**acl\_not\_smtp** 159  
**acl\_not\_smtp\_mime** 159, 398  
**acl\_not\_smtp\_start** 159  
**acl\_smtp\_auth** 159  
**acl\_smtp\_connect** 160  
**acl\_smtp\_data** 160  
**acl\_smtp\_data\_prdr** 160  
**acl\_smtp\_dkim** 160  
**acl\_smtp\_etrn** 160  
**acl\_smtp\_expn** 160  
**acl\_smtp\_helo** 160  
**acl\_smtp\_mail** 160  
**acl\_smtp\_mailauth** 160  
**acl\_smtp\_mime** 160, 398  
**acl\_smtp\_notquit** 161  
**acl\_smtp\_predata** 161  
**acl\_smtp\_quit** 161  
**acl\_smtp\_rcpt** 161  
**acl\_smtp\_starttls** 161  
**acl\_smtp\_vrfy** 161  
**add\_environment** 161  
**address\_data** 207  
**address\_retry\_include\_sender** 291  
**address\_test** 207  
**admin\_groups** 161  
**allow\_commands** 284  
**allow\_defer** 243  
**allow\_domain\_literals** 161  
**allow\_fail** 243  
**allow\_fifo** 263  
**allow\_filter** 243  
**allow\_freeze** 243  
**allow\_localhost** 292  
**allow\_mx\_to\_ip** 162  
**allow\_symlink** 263  
**allow\_utf8\_domains** 162  
**auth\_advertise\_hosts** 162  
**authenticated\_sender** 292  
**authenticated\_sender\_force** 292  
**auto\_thaw** 162  
**av\_scanner** 163, 391

## B

**batch\_id** 260, 263, 280, 284  
**batch\_max** 260, 263, 280, 285  
**bcc** 277  
**bi\_command** 163  
**body\_only** 253  
**bounce\_message\_file** 163  
**bounce\_message\_text** 163  
**bounce\_return\_body** 163  
**bounce\_return\_linesize\_limit** 163  
**bounce\_return\_message** 164  
**bounce\_return\_size\_limit** 164  
**bounce\_sender\_authentication** 164

## C

**callout\_domain\_negative\_expire** 164  
**callout\_domain\_positive\_expire** 164  
**callout\_negative\_expire** 164  
**callout\_positive\_expire** 164  
**callout\_random\_local\_part** 164  
**cannot\_route\_message** 207  
**caseful\_local\_part** 208  
**cc** 278  
**check\_ancestor** 243  
**check\_group** 243, 263  
**check\_local\_user** 208  
**check\_log\_inodes** 165  
**check\_log\_space** 165  
**check\_owner** 244, 264  
**check\_rfc2047\_length** 165  
**check\_secondary\_mx** 223  
**check\_spool\_inodes** 165  
**check\_spool\_space** 165  
**check\_srv** 223  
**check\_string** 264, 285  
**chunking\_advertise\_hosts** 166  
**client\_condition** 318  
**client\_domain** 335  
**client\_ignore\_invalid\_base64** 325  
**client\_name** 328  
**client\_password** 335  
**client\_secret** 328  
**client\_send** 325  
**client\_set\_id** 319  
**client\_username** 335  
**command** 236, 280, 285  
**command\_group** 236  
**command\_timeout** 292  
**command\_user** 236  
**commandline\_checks\_require\_admin** 166  
**condition** 208  
**connect\_timeout** 292  
**connection\_max\_messages** 292  
**create\_directory** 264  
**create\_file** 264  
**current\_directory** 236, 253

## D

**daemon\_smtp\_ports** 166  
**daemon\_startup\_retries** 166  
**daemon\_startup\_sleep** 166  
**data** 244  
**data\_timeout** 293  
**debug\_print** 209, 253

- debug\_store 166
- delay\_after\_cutoff 293, 315
- delay\_warning 166
- delay\_warning\_condition 131, 167
- deliver\_drop\_privilege 167
- deliver\_queue\_load\_max 167
- delivery\_date\_add 253
- delivery\_date\_remove 167, 419
- directory 264
- directory\_file 265
- directory\_mode 265
- directory\_transport 244
- disable\_fsync 167
- disable\_ipv6 151, 168
- disable\_logging 210, 253
- dkim\_canon 293, 481
- dkim\_hash 293, 481
- dkim\_identity 293, 481
- dkim\_private\_key 293, 480
- dkim\_selector 293
- dkim\_sign\_headers 293, 481
- dkim\_strict 293, 481
- dkim\_verify\_signers 168
- dns\_again\_means\_nonexist 168
- dns\_check\_names\_pattern 168
- dns\_csa\_search\_limit 168
- dns\_csa\_use\_reverse 169
- dns\_dnssec\_ok 169
- dns\_ipv4\_lookup 169
- dns\_qualify\_single 293
- dns\_retrans 169
- dns\_retry 169
- dns\_search\_parents 293
- dns\_trust\_aa 169
- dns\_use\_edns0 170
- dnssec\_request\_domains 210, 294
- dnssec\_require\_domains 210, 294
- domains 210
- driver 210, 253, 319
- drop\_cr 170
- dscp 294
- dsn\_advertise\_hosts 170
- dsn\_from 170
- dsn\_lasthop 210

## E

- envelope\_to\_add 253
- envelope\_to\_remove 170, 419
- environment 285
- errors\_copy 170
- errors\_reply\_to 171, 271
- errors\_to 210, 433, 435
- escape\_string 265, 285
- event\_action 171, 254
- exim\_group 171
- exim\_path 171
- exim\_user 171
- expn 211
- extra\_local\_interfaces 172
- extract\_addresses\_remove\_arguments 172

## F

- fail\_defer\_domains 223
- fail\_verify 211

- fail\_verify\_recipient 211
- fail\_verify\_sender 211
- fallback\_hosts 211, 294
- file 244, 265, 278
- file\_expand 278
- file\_format 265
- file\_must\_exist 266
- file\_optional 278
- file\_transport 244
- filter\_prepend\_home 245
- final\_timeout 294
- finduser\_retries 172
- forbid\_blackhole 245
- forbid\_exim\_filter 245
- forbid\_file 245
- forbid\_filter\_dfunc 245
- forbid\_filter\_existstest 245
- forbid\_filter\_logwrite 245
- forbid\_filter\_lookup 245
- forbid\_filter\_perl 245
- forbid\_filter\_readfile 245
- forbid\_filter\_readsocket 246
- forbid\_filter\_reply 246
- forbid\_filter\_run 246
- forbid\_include 246
- forbid\_pipe 246
- forbid\_sieve\_filter 246
- forbid\_smtp\_code 246
- force\_command 285
- freeze\_exec\_fail 285
- freeze\_signal 285
- freeze\_tell 172
- from 278

## G

- gecos\_name 172
- gecos\_pattern 173
- gethostbyname 295
- gnutls\_compat\_mode 173, 295
- group 212, 254

## H

- header\_line\_maxsize 173
- header\_maxsize 173
- headers 278
- headers\_add 212, 254
- headers\_charset 173
- headers\_only 254
- headers\_remove 212, 254
- headers\_rewrite 254
- helo\_accept\_junk\_hosts 173
- helo\_allow\_chars 173
- helo\_data 295
- helo\_lookup\_domains 174
- helo\_try\_verify\_hosts 174
- helo\_verify\_hosts 174
- hide\_child\_in\_errmsg 246
- hold\_domains 174
- home\_directory 255
- host\_all\_ignored 229
- host\_find\_failed 229
- host\_lookup 175
- host\_lookup\_order 175
- host\_reject\_connection 175

- hosts 227, 295
- hosts\_avoid\_esmtp 296
- hosts\_avoid\_pipelining 296
- hosts\_avoid\_tls 296
- hosts\_connection\_nolog 175
- hosts\_max\_try 296
- hosts\_max\_try\_hardlimit 296
- hosts\_nopass\_tls 296
- hosts\_noproxy\_tls 296
- hosts\_override 297
- hosts\_proxy 176
- hosts\_randomize 230, 297
- hosts\_request\_ocsp 297
- hosts\_require\_auth 297
- hosts\_require\_ocsp 297
- hosts\_require\_tls 297
- hosts\_treat\_as\_local 92, 176
- hosts\_try\_auth 297
- hosts\_try\_chunking 298
- hosts\_try\_fastopen 298
- hosts\_try\_prdr 298
- hosts\_verify\_avoid\_tls 296

## I

- ibase\_servers 176
- ignore\_bounce\_errors\_after 10, 176
- ignore\_eaccess 246
- ignore\_enotdir 247
- ignore\_fromline\_hosts 176
- ignore\_fromline\_local 176
- ignore\_quota 280
- ignore\_status 286
- ignore\_target\_hosts 213
- include\_directory 247
- initgroups 213, 251, 255
- interface 298

## K

- keep\_environment 177
- keep\_malformed 177
- keepalive 298

## L

- ldap\_ca\_cert\_dir 177
- ldap\_ca\_cert\_file 177
- ldap\_cert\_file 177
- ldap\_cert\_key 177
- ldap\_cipher\_suite 177
- ldap\_default\_servers 178
- ldap\_require\_cert 178
- ldap\_start\_tls 178
- ldap\_version 178
- lmtpl\_ignore\_quota 299
- local\_from\_check 178
- local\_from\_prefix 178
- local\_from\_suffix 179
- local\_interfaces 179
- local\_part\_prefix 213
- local\_part\_prefix\_optional 214
- local\_part\_suffix 214
- local\_part\_suffix\_optional 214
- local\_parts 214
- local\_scan\_timeout 179

- local\_sender\_retain 179
- localhost\_number 9, 179
- lock\_fcntl\_timeout 266
- lock\_flock\_timeout 266
- lock\_interval 266
- lock\_retries 266
- lockfile\_mode 266
- lockfile\_timeout 266
- log 278
- log\_as\_local 215
- log\_defer\_output 286
- log\_fail\_output 286
- log\_file\_path 180
- log\_output 286
- log\_selector 180
- log\_timezone 180
- lookup\_open\_max 180

## M

- mailbox\_filecount 267
- mailbox\_size 267
- maildir\_format 267
- maildir\_quota\_directory\_regex 267
- maildir\_retries 267
- maildir\_tag 267
- maildir\_use\_size\_file 267
- maildirfolder\_create\_regex 268
- mailstore\_format 268
- mailstore\_prefix 268
- mailstore\_suffix 268
- max\_output 286
- max\_parallel 255
- max\_rcpt 299
- max\_username\_length 180
- mbx\_format 268
- message\_body\_newlines 135, 180
- message\_body\_visible 135, 181
- message\_id\_header\_domain 181
- message\_id\_header\_text 181, 420
- message\_logs 181
- message\_prefix 268, 286
- message\_size\_limit 181, 255
- message\_suffix 269, 287
- mode 269, 278
- mode\_fail\_narrower 269
- modemask 247
- more 215, 218, 229
- move\_frozen\_messages 182
- mua\_wrapper 182, 440
- multi\_domain 299
- mx\_domains 223
- mx\_fail\_domains 224
- mysql\_servers 182

## N

- never\_mail 278
- never\_users 182
- no\_xxx 55
- not\_xxx 55
- notify\_comsat 269

## O

- once 278

once\_file\_size 279  
once\_repeat 279  
one\_time 247  
openssl\_options 182  
optional 227  
oracle\_servers 184  
owners 247  
owngroups 248

## P

pass\_on\_timeout 215  
pass\_router 215  
path 287  
percent\_hack\_domains 184  
perl\_at\_start 147, 184  
perl\_startup 147, 184  
perl\_taintmode 147  
permit\_coredump 287  
pgsql\_servers 184  
pid\_file\_path 185  
pipe\_as\_creator 287  
pipe\_transport 248  
pipelining\_advertise\_hosts 185  
port 227, 299  
prdr\_enable 185, 351  
preserve\_message\_logs 185, 453  
primary\_hostname 92, 185  
print\_topbitchars 185  
process\_log\_path 186  
prod\_requires\_admin 186  
protocol 227, 299  
public\_name 319

## Q

qualify\_domain 186, 248, 417  
qualify\_preserve\_domain 248  
qualify\_recipient 186, 417  
qualify\_single 224  
query 227  
queue\_domains 186  
queue\_list\_requires\_admin 186  
queue\_only 186, 365  
queue\_only\_file 187  
queue\_only\_load 187  
queue\_only\_load\_latch 187  
queue\_only\_override 187  
queue\_run\_in\_order 187  
queue\_run\_max 188  
queue\_smtp\_domains 188  
quota 269  
quota\_directory 270  
quota\_filecount 270  
quota\_is\_inclusive 270  
quota\_size\_regex 270  
quota\_warn\_message 171, 270  
quota\_warn\_threshold 271

## R

rcpt\_include\_affixes 255  
receive\_timeout 188  
received\_header\_text 188  
received\_headers\_max 189  
recipient\_unqualified\_hosts 189

recipients\_max 189  
recipients\_max\_reject 189  
redirect\_router 215  
remote\_max\_parallel 189  
remote\_sort\_domains 190  
repeat\_use 248  
reply\_to 279  
reply\_transport 248  
require\_files 216  
reroute 227  
response\_pattern 228  
restrict\_to\_path 287  
retry\_data\_expire 190, 315  
retry\_include\_ip\_address 299  
retry\_interval\_max 190, 314  
retry\_use\_local\_part 217, 256  
return\_fail\_output 287  
return\_message 279  
return\_output 287  
return\_path 256, 435  
return\_path\_add 256  
return\_path\_remove 190, 420  
return\_size\_limit 190  
rewrite 248  
rewrite\_headers 224  
rfc1413\_hosts 190  
rfc1413\_query\_timeout 191  
route\_data 230  
route\_list 230  
router\_home\_directory 217

## S

same\_domain\_copy\_routing 224, 230  
search\_parents 225  
self 215, 217  
    in *dnslookup* router 222  
    in *ipliteral* router 226  
    in *manualroute* router 232  
    value of host name 140  
sender\_unqualified\_hosts 191  
senders 218  
serialize\_hosts 300  
server\_advertise\_condition 319  
server\_channelbinding 332  
server\_condition 319, 323  
server\_debug\_print 319  
server\_hostname 329, 332, 334  
server\_keytab 334  
server\_mail\_auth\_condition 320  
server\_mech 329, 332  
server\_param1 337  
server\_param2 337  
server\_param3 337  
server\_password 332, 335  
server\_prompts 323  
server\_realm 329, 332  
server\_scam\_iter 333  
server\_scam\_salt 333  
server\_secret 327  
server\_service 330, 333, 334  
server\_set\_id 319  
server\_socket 331  
set\_environment 191  
shadow\_condition 256  
shadow\_transport 256

- sieve\_subaddress 248
- sieve\_useraddress 249
- sieve\_vacation\_directory 249
- size\_addition 300
- skip\_syntax\_errors 249
- slow\_lookup\_log 191
- smtp\_accept\_keepalive 191
- smtp\_accept\_max 191
- smtp\_accept\_max\_nonmail 191
- smtp\_accept\_max\_nonmail\_hosts 192
- smtp\_accept\_max\_per\_connection 192
- smtp\_accept\_max\_per\_host 192
- smtp\_accept\_queue 192
- smtp\_accept\_queue\_per\_connection 192
- smtp\_accept\_reserve 193
- smtp\_active\_hostname 193
- smtp\_banner 193, 350, 361
- smtp\_check\_spool\_space 193
- smtp\_connect\_backlog 194
- smtp\_enforce\_sync 194
- smtp\_etrn\_command 132, 194, 429
- smtp\_etrn\_serialize 194
- smtp\_load\_reserve 194
- smtp\_max\_synprot\_errors 195
- smtp\_max\_unknown\_commands 195
- smtp\_ratelimit\_\* 378
- smtp\_ratelimit\_hosts 195
- smtp\_ratelimit\_mail 195
- smtp\_ratelimit\_rcpt 196
- smtp\_receive\_timeout 196
- smtp\_reserve\_hosts 196
- smtp\_return\_error\_details 196
- smtpUTF8\_advertise\_hosts 196
- socket 280
- socks\_proxy 300
- spamd\_address 196, 396
- split\_spool\_directory 197
- spool\_directory 197
- spool\_wireformat 197
- sqlite\_lock\_timeout 198
- srv\_fail\_domains 225
- strict\_acl\_vars 198, 357
- strip\_excess\_angle\_brackets 198
- strip\_trailing\_dot 198
- subject 279
- syntax\_errors\_text 250
- syntax\_errors\_to 250
- syslog\_duplication 198
- syslog\_facility 198
- syslog\_pid 198
- syslog\_processname 198
- syslog\_timestamp 199
- system\_filter 199
- system\_filter\_directory\_transport 199
- system\_filter\_file\_transport 199
- system\_filter\_group 199
- system\_filter\_pipe\_transport 199
- system\_filter\_reply\_transport 199
- system\_filter\_user 199

## T

- tcp\_nodelay 200
- temp\_errors 288
- text 279
- timeout 228, 236, 280, 288

- timeout\_defer 288
- timeout\_frozen\_after 10, 200
- timezone 200
- tls\_advertise\_hosts 200
- tls\_certificate 300
- tls\_crl 201, 300
- tls\_dh\_max\_bits 201
- tls\_dh\_min\_bits 300
- tls\_dhparam 201
- tls\_eccurve 202
- tls\_in\_sni 345
- tls\_ocsp\_file 202
- tls\_on\_connect\_ports 202
- tls\_privatekey 301
- tls\_remember\_esmtp 203
- tls\_require\_ciphers 203, 301
  - GnuTLS 341
  - OpenSSL 340
- tls\_sni 301
- tls\_tempfail\_tryclear 301
- tls\_try\_verify\_hosts 203, 301
- tls\_verify\_cert\_hostnames 301
- tls\_verify\_certificates 203, 302
- tls\_verify\_hosts 204, 302
- to 279
- translate\_ip\_address 218
- transport 219
- transport\_current\_directory 219
- transport\_filter 257
- transport\_filter\_timeout 258
- transport\_home\_directory 219
- trusted\_groups 204
- trusted\_users 204

## U

- umask 288
- unknown\_login 204
- unknown\_username 204
- unseen 12, 212, 219, 422
- unset 200, 203, 293, 480
- untrusted\_set\_sender 204
- use\_bsmtplib 271, 288
- use\_classresources 288
- use\_crlf 271, 288
- use\_fcntl\_lock 271
- use\_flock\_lock 272
- use\_lockfile 272
- use\_mbx\_lock 272
- use\_shell 289
- user 220, 258
- uucp\_from\_pattern 205, 418
- uucp\_from\_sender 205, 418

## V

- verify 220
- verify\_only 220
- verify\_recipient 220
- verify\_sender 220

## W

- warn\_message\_file 205
- widen\_domains 225
- write\_rejectlog 206

# Variables index

## Symbols

\$1, \$2, etc. *see numerical variables*

\$acl\_smtp\_notquit 352

\$acl\_verify\_message 128, 242, 356, 360, 361, 382

\$address\_data 128, 141, 207, 211, 237, 372

\$address\_file 129, 199, 244, 262

\$address\_pipe 129, 199, 248, 282

\$auth1 333, 334

\$auth1, \$auth2, etc 129, 323, 333, 334

\$auth2 333, 334

\$auth3 333

\$authenticated\_fail\_id 129

\$authenticated\_id 129, 319, 323, 327, 419, 421

\$authenticated\_sender 129, 320

\$authentication\_failed 129

\$bheader\_ 106

\$body\_linecount 130, 475

\$body\_zerocount 130

\$bounce\_recipient 130, 431

\$bounce\_return\_size\_limit 130, 431

\$caller\_gid 130, 137

\$caller\_uid 130, 137, 204

\$callout\_address 130, 395, 397

\$compile\_number 130

\$config\_dir 130

\$config\_file 130

\$dkim\_signers 131

\$dnslist\_domain 131, 375

\$dnslist\_matched 131, 375

\$dnslist\_text 131, 375

\$dnslist\_value 131, 375

\$domain 131, 137, 167, 171, 194, 235, 252, 260, 299, 305, 306, 307, 308, 354, 370, 412, 415, 429

\$domain\_data 132, 210, 369

\$exim\_gid 132

\$exim\_path 132

\$exim\_uid 132

\$exim\_version 132

\$header\_ 106

\$headers\_added 132

\$home 13, 132, 208, 217, 255

\$host 132, 229, 233, 235, 258, 291, 298, 300, 301, 302, 322, 328, 345, 430

\$host\_address 133, 213, 218, 258, 291, 298, 300, 301, 302, 322, 328, 345

\$host\_data 133, 370

\$host\_lookup\_deferred 133, 142

\$host\_lookup\_failed 133, 142, 175

\$host\_port 133

\$initial\_cwd 133

\$inode 133, 265

\$interface\_address 133

\$interface\_port 133

\$item 106, 109, 112, 123, 133

\$ldap\_dn 133

\$load\_average 134

\$local\_part 13, 134, 137, 171, 208, 213, 239, 252, 260, 262, 282, 305, 306, 307, 308, 354, 362, 412, 415, 435

\$local\_part\_data 134, 214, 370

\$local\_part\_prefix 13, 134, 213

\$local\_part\_suffix 13, 67, 134, 437

\$local\_scan\_data 135, 403

\$local\_user\_gid 13, 135

\$local\_user\_uid 13, 135

\$localhost\_number 135, 179

\$log\_inodes 135, 165

\$log\_space 135, 165

\$lookup\_dnssec\_authenticated 135

\$mailstore\_basename 135

\$malware\_name 135, 395

\$max\_received\_linelength 135

\$message\_age 135

\$message\_body 135, 180, 181

\$message\_body\_end 136, 180, 181

\$message\_body\_size 136

\$message\_exim\_id 136

\$message\_headers 136

\$message\_headers\_raw 136

\$message\_linecount 136

\$message\_size 136, 275, 354

\$original\_domain 137, 252

\$original\_local\_part 137, 208

\$originator\_gid 137

\$originator\_uid 137

\$parent\_domain 137

\$parent\_local\_part 137, 208

\$pid 137

\$pipe\_addresses 137, 257, 260, 282, 283, 289

\$port 299

\$primary\_hostname 137, 185, 193

\$qualify\_domain 31, 129, 138, 385, 419, 421

\$qualify\_recipient 138, 248

\$queue\_name 138

\$rcpt\_count 138, 354

\$rcpt\_defer\_count 138

\$rcpt\_fail\_count 138

\$received\_count 138

\$received\_for 138

\$received\_ip\_address 138

\$received\_port 139

\$received\_protocol 46, 139, 321

\$received\_time 139

\$recipient\_data 139

\$recipient\_verify\_failure 139, 382

\$recipients 139, 413

\$recipients\_count 140, 354

\$regex\_match\_string 140

\$reply\_address 140

\$return\_path 140, 256

\$return\_size\_limit 140

\$rheader\_ 106

\$router\_name 140

\$run\_in\_acl 112

\$runrc 112, 140

\$self\_hostname 140, 218

\$sender\_address 140, 277, 304, 320, 354

\$sender\_address\_data 141, 207, 372

\$sender\_address\_domain 141, 370

\$sender\_address\_local\_part 141

\$sender\_data 141

\$sender\_fullhost 141, 451

\$sender\_helo\_dnssec 141

\$sender\_helo\_name 141



*\$sender\_host\_address* 141, 354, 409  
*\$sender\_host\_authenticated* 141, 321  
*\$sender\_host\_dnssec* 141  
*\$sender\_host\_name* 142, 175  
*\$sender\_host\_port* 142  
*\$sender\_ident* 142, 205  
*\$sender\_rcvhost* 142, 451  
*\$sender\_verify\_failure* 143, 382  
*\$sending\_ip\_address* 143  
*\$sending\_port* 143  
*\$smtp\_active\_hostname* 143, 193  
*\$smtp\_command* 143, 354  
*\$smtp\_command\_argument* 143, 354  
*\$smtp\_command\_history* 143  
*\$smtp\_count\_at\_connection\_start* 143  
*\$smtp\_notquit\_reason* 352  
*\$spool\_directory* 144  
*\$spool\_inodes* 144, 165  
*\$spool\_space* 144, 165  
*\$thisaddress* 144  
*\$tls\_bits* 291  
*\$tls\_cipher* 145, 291  
*\$tls\_in\_bits* 144  
*\$tls\_in\_certificate\_verified* 145  
*\$tls\_in\_cipher* 145, 162, 321, 343  
*\$tls\_in\_ocsp* 145  
*\$tls\_in\_ourcert* 144  
*\$tls\_in\_peercert* 144  
*\$tls\_in\_peerdn* 145, 339, 343  
*\$tls\_in\_sni* 146, 345  
*\$tls\_out\_bits* 144, 345  
*\$tls\_out\_certificate\_verified* 145  
*\$tls\_out\_cipher* 145, 345  
*\$tls\_out\_ocsp* 145  
*\$tls\_out\_ourcert* 144  
*\$tls\_out\_peercert* 145  
*\$tls\_out\_peerdn* 145, 339, 345  
*\$tls\_out\_sni* 146, 345  
*\$tls\_peerdn* 145, 291  
*\$tls\_sni* 146, 291, 301  
*\$tod\_bsdinbox* 146  
*\$tod\_epoch* 146  
*\$tod\_epoch\_l* 146  
*\$tod\_full* 146  
*\$tod\_log* 146, 180  
*\$tod\_logfile* 146  
*\$tod\_zone* 146, 180  
*\$tod\_zulu* 146  
*\$transport\_name* 146  
*\$value* 105, 109, 112, 146, 231  
*\$verify\_mode* 146  
*\$version\_number* 146  
*\$warn\_message\_delay* 146, 432  
*\$warn\_message\_recipients* 146, 432

## T

**tls\_certificate** 346  
**tls\_crl** 346  
**tls\_ocsp\_file** 346  
**tls\_privatekey** 346  
**tls\_verify\_certificates** 346

# Concept index

## Symbols

- `.ifdef` 55
- `.include` in configuration file 53
- `.include_if_exists` in configuration file 53
- `.so` building 21
- `@` in a domain list 60, 92
- `@` in a host list 94
- `@@` with single-key lookup 100
- `@[]` in a domain list 92
- `@[]` in a host list 95
- `@mx_any` 92
- `@mx_primary` 92
- `@mx_secondary` 92
- `+caseful` 101, 125
- `+defer_unknown` 373
- `+exclude_unknown` 373
- `+ignore_defer` 98
- `+ignore_unknown` 97
- `+include_defer` 98
- `+include_unknown` 97, 373
- `*@` with single-key lookup 76, 100
- `/dev/null` 241
- `/etc/aliases` 24
- `/etc/mail/mailer.conf` 26
- `/etc/passwd` 76, 208
  - multiple reading of 172
- `/etc/userdbshadow.dat` 73

## Digits

- 4xx responses
  - count of 138
  - retry rules for 311
  - retrying after 291
  - to STARTTLS 301
- 8-bit characters 29, 159, 185
- 8BITMIME 159, 450

## A

- A+
  - in *dnsdb* lookup 80
- abandoning mail 42, 241
- accept** ACL verb 355
- accept* router 221
- access control lists (ACLs)
  - arguments 368
  - at start of non-SMTP message 159
  - case of local part in 362
  - certificate verification 370
  - conditions; list of 368
  - conditions; processing 357
  - customized test 369
  - cutthrough routing 362
  - data for message ACL 354
  - data for non-message ACL 354
  - default configuration 63
  - description 349–390
  - enabling debug logging 363
  - finding which to use 352
  - for non-SMTP messages 159
  - format of 354
  - indirect 368

access control lists (ACLs) (*continued*)

- introduction 8
- modifiers; list of 358
- modifiers; processing 357
- nested 368
- on SMTP connection 160
- options for specifying 349
- per-user data processing 160
- PRDR-related 160
- relay control 389
- return codes 353
- rewriting addresses in 304
- scanning for spam 370
- scanning for viruses 370
- setting DSCP value 363
- setting up for SMTP commands 159
- testing a DNS list 369, 372
- testing a local part 370
- testing a recipient 370
- testing a recipient domain 369
- testing a sender 370
- testing a sender domain 370
- testing a TLS certificate 370
- testing by regex matching 370
- testing for authentication 369
- testing for encryption 369
- testing the client host 369
- testing, customized 369
- unset options 353
- variables 356
- variables, handling unset 198
- verbs, definition of 354
- verifying header names only ASCII 371
- verifying header syntax 371
- verifying HELO/EHLO 371
- verifying host reverse lookup 372
- verifying recipient 372
- verifying sender 372
- verifying sender in the header 371
- virus scanning 370

## acl

- call from expansion 103
- expansion condition 121

**acl** ACL condition 368

**add\_header** ACL modifier 366

adding drivers 491

additional groups 213, 255

address

- constructed 422
- copying routing 224, 230
- qualification 186, 417
- rewriting *see* *rewriting*
- sender 40
- source-routed 184
- testing 36, 211
- verification 36
- without domain 4

address duplicate, discarding 13, 242

**address** expansion item 114

address list

- `@@` lookup type 100
- case forcing 101

- address list (*continued*)
  - empty item 99
  - in a rewriting pattern 306
  - in expansion condition 125
  - local part starting with ! 99
  - lookup for complete address 99
  - patterns 99
  - regular expression in 99
  - split local part and domain 100
- address qualification, suppressing 33
- address redirection
  - broken files 249
  - disabling rewriting 248
  - domain; preserving 248
  - errors 242
  - included external list 241
  - local part without domain 239
  - non-filter list items 239
  - one-time expansion 247
  - redirect* router 238
  - repeated for each delivery attempt 242
  - to black hole 241
  - to file 240
  - to local mailbox 239
  - to pipe 240
  - while verifying 238, 387
- addresses** expansion item 114
- admin user 161, 464, 473
  - definition of 28
- alias file
  - backslash in 239
  - broken 249
  - building 28, 32
  - exception to default 242
  - in a *redirect* router 238
  - one-time expansion 247
  - ownership 247
  - per-domain default 76
- alias for host 96
- alternate configuration file 37
- “and” expansion condition 128
- angle brackets, excess 198
- appendfile* transport 262–276
- appending to a file 272
- asterisk
  - after IP address 425, 447
  - in address list 100
  - in domain list 93
  - in host list 94, 96
  - in lookup type 93
  - in search type 77
- Athena 7
- AUTH
  - ACL for 159, 349
  - advertising 162
  - advertising when encrypted 162
  - configuration 52, 69
  - description of 317
  - in *plaintext* authenticator 323
  - logging 446
  - on bounce message 164
  - on MAIL command 129, 160, 320, 322
  - testing a server 321
  - with PAM 126
- authenticated** ACL condition 369

- authentication 317–322
  - ACL checking 369
  - advertising 162
  - ANONYMOUS 332, 333
  - bounce message 164
  - Client Certificate 337
  - CRAM-MD5 332
  - CRAM-MD5 mechanism 327
  - DIGEST-MD5 332
  - EXTERNAL 332, 333
  - fail 129
  - failure 129
  - generic options 318
  - GNU SASL 332
  - GSSAPI 333, 334
  - id 129
  - id, specifying for local message 45
  - Kerberos 334
  - logging 446
  - LOGIN 332
  - LOGIN mechanism 324
  - Microsoft Secure Password 335
  - name, specifying for local message 45
  - NTLM 335
  - on an Exim client 322
  - on an Exim server 320
  - optional in client 297
  - PLAIN 332
  - PLAIN mechanism 323
  - required by client 297
  - SASL 332
  - SCRAM-SHA-1 332
  - sender 129
  - sender, specifying for local message 45
  - sender; authenticated 320
  - testing a server 321
  - to proxy 486
  - X509 337
- authenticators
  - cram\_md5* 327–328
  - cyrus\_sasl* 329–330
  - dovecot* 331
  - gsasl* 332
  - heimdal\_gssapi* 334
  - plaintext* 323–326
  - spa* 335–336
  - tls* 337
- autoreply* transport 277–279
  - for system filter 199

## B

- background delivery 43
- backlog of connections 194
- backslash in alias file 239
- bang paths
  - not handled by Exim 3
  - rewriting 309
- banner for SMTP 193
- base32** expansion item 115
- base32d** expansion item 115
- base36 9
- base62 9, 265
- base62** expansion item 115
- base62d** expansion item 115

- base64 decoding
  - in string expansion 115
- base64 encoding
  - conversion from hex 117
  - creating authentication test data 321
  - functions for *local\_scan()* use 408
  - in encrypted password 122
  - in header lines 106
  - in *plaintext* authenticator 323
  - in string expansion 115
- base64** expansion item 115
- base64d** expansion item 115
- batched local delivery 260
- batched SMTP input 35, 430
- batched SMTP output 429
- batched SMTP output example 234
- BATV, verifying 388
- bcc recipients, verifying none 371
- Bcc*: header line 49, 419
- BDAT
  - SMTP command 298, 351
- Berkeley DB library 18
  - file format 73
- BIN\_DIRECTORY 24
- binary zero
  - in authentication data 321
  - in header line 106
  - in lookup key 73, 74, 459
  - in message body 130, 135
  - in *plaintext* authenticator 323
  - in RFC 2047 decoding 410
- bind IP address 298
- black hole 241
- black list (DNS) 131, 369, 372, 451
- body of message
  - binary zero count 130
  - definition of 4
  - expansion variable 135, 136
  - line count 130
  - size 136
  - transporting 253
  - visible size 181
- books about Exim 1
- bool** expansion condition 121
- bool\_lax** expansion condition 121
- boolean configuration values 55
- Bounce Address Tag Validation *see* BATV
- bounce message
  - copy to other address 170
  - customizing 163, 431
  - definition of 4
  - discarding 176
  - failure to deliver 16
  - generating 39
  - including body 163
  - including original 163
  - line length limit 163
  - recipient of 15
  - redirection details; suppressing 246
  - Reply-to*: in 171
  - sender authentication 164
  - size limit 164
  - when generated 15
- bounce messages
  - From*: line, specifying 170

- bounce messages (*continued*)
  - success 170
- broken alias or forward files 249
- BSD, DBM library for 17
- bug reports 3
- Bugzilla 2
- build directory 21
- build-time options, overriding 21
- building alias file 32
- building DBM files 459
- building Exim 17–24
  - architecture type 22
  - multiple OS/architectures 17
  - operating system type 22
  - OS-specific C header files 23
  - overriding default settings 22
  - pre-building configuration 18
- building Eximon 23

## C

- caching
  - callout 386
  - callout timeouts 164
  - lookup data 78
  - named lists 91
  - of dns lookup 80, 373
- caching callout, suppressing 384
- callout
  - additional parameters for 384
  - cache, description of 386
  - cache, suppressing 384
  - caching timeouts 164
  - connection timeout, specifying 384
  - defer, action on 384
  - full postmaster check 384
  - overall timeout, specifying 384
  - postmaster; checking 384
  - “random” check 385
  - sender for recipient check 385
  - sender when verifying header 384
  - timeout, specifying 384
  - verification 383
- carriage return 271, 288, 417, 424, 427
- case forcing in address lists 101
- case forcing in strings 117, 120
- case of local parts 13, 101, 208, 362, 423
- case sensitivity
  - in (n)wildsearch lookup 74
  - in lsearch lookup 74
- Cc*: header line 49
- cdb
  - acknowledgment 6
  - description of 73
  - including support for 22
- certextract**
  - certificate fields 103
- certificate
  - base64 of DER 115
  - client, location of 300
  - directory for LDAP 177
  - extracting fields 103, 145
  - file for LDAP 177
  - fingerprint 118, 119
  - key for LDAP 177
  - references to discussion 347

- certificate (*continued*)
  - revocation list 343
  - revocation list for client 300
  - revocation list for server 201
  - self-signed 347
  - server, location of 200
  - variables 144
  - verification of client 203, 204, 343, 370
  - verification of server 301, 302
- Certificate-based authentication 337
- change log 2
- checking access 459
- checking disk space 165, 193
- CHUNKING
  - advertisement 166
  - BDAT command 351
  - enabling, in client 298
- CIDR notation 95, 118
- cipher
  - logging 343, 344
  - requiring specific 203, 301
- class resources (BSD) 288
- Client SMTP Authorization *see* CSA
- client, non-queueing 440
- command line
  - addresses with **-t** 172
  - options 28–50
- common option syntax 55
- compiler
  - requirements 17
  - version 17
- concurrent deliveries 251
- condition** ACL condition 369
- configuration file
  - alternate 37, 51
  - common option syntax 55
  - conditional skips 55
  - default “walk through” 60–70
  - editing 23
  - errors in 51
  - format of 52
  - general description 51
  - including other files 53
  - leading white space in 52
  - macros 53
  - main section 153–206
  - ownership 51
  - retry section 310–316
  - trailing white space in 52
- configuration for building Exim 18
- configuration options
  - extracting 33
  - hiding value of 33, 55
- CONFIGURE\_FILE 24, 37, 51
- CONFIGURE\_GROUP 51
- CONFIGURE\_OWNER 51
- connection backlog 194
- constructed address 422
- content scanning
  - at ACL time 391–401
  - AV scanner failure 130
  - for spam 395
  - for viruses 391
  - MIME parts 398
  - per user 257
- content scanning (*continued*)
  - specifying at build time 19
  - the **malware** condition 391
  - with regular expressions 401
- continue** ACL modifier 358
- control** ACL modifier 358, 362
- control of incoming mail 349
- copy of bounce message 170
- copy of message (**unseen** option) 219
- correction of invalid utf-8 sequences in strings 120
- Courier 73
- CR character *see* carriage return
- cram\_md5 authenticator 327–328
- CRAM-MD5 authentication mechanism 327
- creating directories 262
- CRL *see* certificate revocation list
- crypt() 122
- crypt16() 122
- crypteq** expansion condition 121
- CSA
  - in *dnsdb* lookup 80
  - verifying 387
- CSA verification 371
- current directory for local transport 219, 252
- customizing
  - ACL condition 369
  - batching condition 260
  - bounce message 163, 431
  - “cannot route” message 207
  - failure message 241
  - input scan using C function 402
  - precondition 14
  - Received:* header 188
  - SMTP banner 193
  - warning message 205, 432
- cutthrough
  - logging 447
  - requesting 362
- cycling logs 443, 458
- Cygwin 9
- Cyrus 6, 126, 127, 286, 289, 292
  - SASL library 329
- cyrus\_sasl* authenticator 329–330

## D

- daemon 29, 37, 426
  - listening IP addresses 149
  - pid file path 185
  - process id (pid) 29, 33, 46
  - restarting 30, 471
  - starting 149
  - TCP\_NODELAY on sockets 200
- daemon startup, retrying 166
- Darwin 9
- DATA
  - ACL for 160
  - ACLs for 349, 350
  - PRDR ACL for 160
- database
  - lookups 72–88
  - updating in ACL 358
- Date:* header line 419
- DBM
  - building dbm files 459
  - lookup type 73

- DBM libraries
  - configuration for building 18, 23
  - discussion of 17
- dbmz lookup type 73
- dbmz lookup type 73
- DCSP
  - outbound 294
- Debian
  - information sources 1
  - mailing list for 2
- debugging
  - bh** option 31
  - d** option 38
  - N** option 43
  - enabling from an ACL 363
  - from embedded Perl 148
  - list of selectors 38
  - memory corruption 166
  - suppressing delivery 43
- decode** ACL condition 369
- def** expansion condition 122
- default
  - ACLs 63
  - configuration file “walk through” 60
  - in single-key lookups 76
  - retry rule 69
  - routers 66
  - transports 68
- defer** ACL verb 355
- defer** in system filter 413
- defer, fake 364
- deferred delivery, forcing 241
- delay** ACL modifier 359
- delay warning, specifying 166
- delayed delivery, logging 450
- delivery
  - abandoning further attempts 42
  - blackhole 241
  - by external agent 289
  - cancelling all 42
  - cancelling by address 42
  - cutthrough; logging 447
  - deferral 15
  - delaying certain domains 174
  - discard 241
  - discarded; logging 447
  - failure report *see bounce message*
  - failure; logging 447
  - fake; logging 448
  - first 123
  - forcing attempt 41
  - forcing deferral 241
  - forcing failure 241, 467
  - forcing in queue run 48
  - from given sender 49
  - in detail 14
  - in the background 43
  - in the foreground 44
  - log line format 215
  - manually started – not forced 42
  - maximum number of 190
  - parallelism for remote 189
  - parallelism for transport 255
  - permanent failure 15
  - problems with 26

- delivery (*continued*)
  - procmail* 289
  - retry in remote transports 15
  - retry mechanism 15
  - sorting remote 190
  - suppressing immediate 44
  - temporary failure 15
  - to file; forbidding 245
  - to given domain 49
  - to pipe; forbidding 246
  - to single file 274
  - unprivileged 167
- delivery failure, long-term 315
- Delivery Status Notification
  - success 170, 210, 238
- Delivery-date*: header line 167, 253, 419
- deny** ACL verb 355
- design philosophy 8
- dialup *see intermittently connected hosts*
- directories, multiple 197
- directory creation 262, 264, 272, 274
- disable DKIM verify 363
- discard** ACL verb 355
- discarded messages 447
- discarding bounce message 176
- disk space, checking 165, 193
- distribution
  - ftp site 3
  - public key 3
  - signing details 3
- DKIM 480
  - ACL for 160, 349
  - controlling calls to the ACL 168
  - disable verify 363
  - signing 480
  - verification 481
- dlfunc** 104
- DNS
  - as a lookup type 75, 79
  - DNSSEC 135, 141, 169, 174, 210, 294
  - EDNS0 170
  - IPv6 lookup for AAAA records 169
  - OpenBSD 170
  - pre-check of name syntax 168
  - qualifying single-component names 224
  - resolver options 169, 170, 224, 225
  - reverse lookup 142, 174, 477
  - timeout 80, 169
  - “try again” response; overriding 168
  - TTL 80, 373
- dns
  - logging slow lookups 191
- DNS list
  - data returned from 374
  - in ACL 369, 372
  - information from merged 377
  - IPv6 usage 378
  - keyed by domain name 373
  - keyed by explicit IP address 373
  - logging defer 451
  - matching specific returned data 375
  - multiple keys for 374
  - variables set from 375
- DNS resolver, debugging output 39
- dnsdb lookup 79

- dnsdb modifiers 79
- dnslists** ACL condition 369
- dnslookup* router 222–225
  - declines 222
- DNSSEC
  - dns lookup 80
  - MX lookup 210, 294
- dnssec
  - logging 451
- doc/ChangeLog* 2
- doc/NewStuff* 2
- doc/spec.txt* 2
- documentation 1
  - available formats 3
- domain
  - ACL checking 369
  - definition of 5
  - delaying delivery 174
  - delivery to 49
  - extraction 115
  - for qualifying addresses 186
  - in redirection; preserving 248
  - manually routing 229
  - partial; widening 225
  - specifying non-immediate delivery 186
  - UTF-8 characters in 162
  - virtual 436
- domain list
  - asterisk in 93
  - in expansion condition 125
  - matching by lookup 93
  - matching “ends with” 93
  - matching literal domain name 93
  - matching local IP interfaces 92
  - matching MX pointers to local host 92
  - matching primary host name 92
  - matching regular expression 93
  - patterns for 92
- domain literal 92, 161
  - default router 66
  - recognizing format 61
  - routing 226
- domainless addresses 4
- domains** ACL condition 369
- dot
  - in incoming non-SMTP message 40, 45
  - in local part 423
  - trailing on domain 198
- dovecot* authenticator 331
- drivers
  - adding new 491
  - configuration format 58
  - definition of 11
  - instance definition 11
- drop** ACL verb 355
- DSCP
  - inbound 363
  - values 32
- dsearch lookup type 73
- DSN
  - success 170, 210, 238
- duplicate addresses 13, 36, 212, 242
- dynamic modules 21

- E**
- EACCES 246
- EAI 120, 487
- EHLO 424, 446
  - accepting junk data 173
  - ACL for 160, 349, 350
  - argument, setting 295
  - avoiding use of 296
  - forcing reverse lookup 174
  - invalid data 427
  - underscores in 173
  - verifying 371
  - verifying, mandatory 174
  - verifying, optional 174
- empty item in hosts list 94
- encrypted** ACL condition 369
- encrypted strings, comparing 121
- encryption
  - checking in an ACL 369
  - including support for 19
  - on SMTP connection 200, 338–348
- endpass** ACL modifier 359
- ENOTDIR 247
- envelope sender 28, 31, 40, 178, 204, 210, 256, 271, 288, 414, 418, 435
  - rewriting at transport time 305
- envelope, definition of 5
- Envelope-to:* header line 170, 253, 260, 419
- environment 34, 191
  - local transports 251–252
  - pipe* transport 284, 285
  - set values 161
  - values from 105, 177, 200
- eq** expansion condition 123
- eqi** expansion condition 123
- error
  - ignored 449
  - in configuration file 51
  - in outgoing SMTP 425
  - reporting 44, 45
  - skipping bad syntax 249
- escape characters in quoted strings 56
- escape** expansion item 115
- escape8bit** expansion item 116
- ESMTP, avoiding use of 296
- ETRN
  - ACL for 160, 349
  - command to be run 194
  - logging 451
  - processing 428
  - serializing 194
  - value of *\$domain* 132
- eval** expansion item 116
- events 171, 254, 489
- exec failure 285
- exicyclog* 443, 458
- exigrep* 457
- Exim arguments, logging 450
- Exim binary, path name 171
- Exim group 171
- Exim monitor
  - acknowledgment 7
  - description 464–468
  - window size 465
- Exim user 171

- exim\_checkaccess* 459
- exim\_dbmbuild* 459
- exim\_dumpdb* 461
- exim\_fixdb* 462
- EXIM\_GROUP 51
- exim\_lock* 462
- exim\_monitor/EDITIME* 19, 464
- exim\_tidydb* 461
- EXIM\_USER 51
- eximon* 464
- eximstats* 458
- exinext* 460
- exipick* 457
- exiqgrep* 455
- exiqsumm* 456
- exiscan* *see content scanning*
- exists**, expansion condition 123
- exiwhat* 186, 455
- expansion
  - “and” of conditions 128
  - arithmetic expression 116
  - base64 decoding 115
  - base64 encoding 115
  - boolean parsing 121
  - calling an acl 103, 121
  - calling Perl from 110
  - case forcing 117, 120
  - character translation 114
  - checking for empty variable 122
  - checking header line existence 122
  - combining conditions 127
  - conditional 108
  - conditions 120–128
  - conversion to base 32 115
  - conversion to base 62 115
  - definition of 57
  - domain extraction 115
  - encrypted comparison 121
  - escape sequences 102
  - escaping 8-bit characters 116
  - escaping non-printing characters 115
  - expression evaluation 116
  - extracting certificate fields 103
  - extracting list elements by number 108
  - extracting substrings by key 105
  - extracting substrings by number 105
  - extracting value from environment 105
  - file existence test 123
  - first delivery test 123
  - forall** condition 123
  - forany** condition 123
  - forced failure 103
  - header insertion 106
  - hex to base64 117
  - hmac hashing 107
  - imap folder 108
  - including literal text 102
  - inserting an entire file 111
  - inserting from a socket 111
  - IP address 119
  - IP address masking 118
  - LDAP authentication test 124
  - list creation 109
  - list item count 117
  - list sorting 113

- expansion (*continued*)
  - local part extraction 118
  - lookup in 109
  - MD5 hash 118
  - named list 117
  - negating a condition 121
  - non-expandable substrings 102
  - numeric comparison 121
  - numeric hash 110, 118
  - of lists 89
  - of strings 102–146
  - operators 103, 114
  - “or” of conditions 127
  - PAM authentication test 126
  - pwcheck* authentication test 126
  - queue runner test 127
  - quoting 118
  - Radius authentication 127
  - re-expansion of substring 116
  - reducing a list to a scalar 112
  - regular expression comparison 124
  - RFC 2047 119
  - RFC 2822 address handling 114
  - running a command 112
  - saslauthd* authentication test 127
  - selecting from list by condition 106
  - SHA-1 hashing 119
  - SHA-256 hashing 119
  - SHA3 hashing 120
  - statting a file 120, 245
  - string comparison 123, 124
  - string length 120
  - string substitution 113
  - string truncation 108, 117
  - substring expansion 120
  - substring extraction 113
  - testing 30, 43, 102
  - textual hash 106, 117
  - UTF-8 120
  - UTF-8 conversion 117
  - utf-8 forcing 120
  - variables 103
  - variables, list of 128
  - variables, set from DNS list 375
- EXPAN
  - ACL for 160, 349
  - error text, display of 241
  - processing 428
  - router skipping 211
  - with **verify\_only** 220
- external local delivery 289
- external transports 4
- extract**
  - substrings by key 105
  - substrings by number 105
- EXTRALIBS 23

## F

- fail**
  - in system filter 413
  - log line; reducing 413
- failing delivery
  - forcing 241
  - from filter 243
- failover *see fallback*



- failure of exec 285
- fake defer 364
- fake rejection 364
- fallback
  - hosts specified on router 211
  - hosts specified on transport 292, 294
  - randomized hosts 297
- fallover *see fallback*
- FAQ 2
- fast open, TCP
  - enabling, in client 298
- fifo (named pipe) 263
- file
  - appending 272
  - existence test 123
  - extracting characteristics 120
  - how a message is held 10
  - in redirection list 240
  - inserting into expansion 111
  - journal 11
  - locking 268, 272, 273
  - lookups 72–88, 109
  - mailbox; checking existing format 265
  - MBX format 268
  - requiring for router 216
  - too many open 180
  - transport for system filter 199
- filter
  - enabling use of 243
  - header lines; adding/removing 414
  - introduction 8
  - locking out certain features 245, 246
  - Sieve *see Sieve filter*
  - specifying in redirection data 239
  - system filter 199, 412–415
  - testing 30
  - transport filter 132, 137, 257, 283, 300, 424
  - user for processing 220
- filtering all mail 412–415
- first delivery 123
- first\_delivery** expansion condition 123
- fixed point configuration values 56
- force command 285
- forcing delivery 41
- foreground delivery 44
- format
  - boolean 55
  - configuration file 52
  - fixed point 56
  - group name 57
  - integer 56
  - list item in configuration 57
  - message 32
  - octal integer 56
  - of message id 9
  - spool files 475–479
  - string 56
  - time interval 56
  - user name 57
- forward file
  - broken 249
  - one-time expansion 247
  - ownership 247
  - testing 30
- FreeBSD, MTA indirection 26

- freeze** in system filter 413
- freezing messages 42, 413
  - allowing in filter 243
  - sending a message when freezing 172
- “From” line 28, 31, 32, 40, 176, 205, 264, 268, 274, 286, 418
- from\_utf8** expansion item 117
- From:* header line 28, 419
  - disabling checking of 178
  - in bounces 170
- frozen messages
  - display 466
  - forcing delivery 41, 48, 49
  - forcing in ACL 364
  - in queue listing 34
  - logging skipping 452
  - manual thaw; testing in filter 412
  - moving 182
  - spool data 476
  - thawing 10, 43, 467
  - timing out 200
- fsync()*, disabling 167
- FTP site 2, 3

## G

- gdbm* DBM library 18
- ge** expansion condition 123
- “gecos” field, parsing 172
- gei** expansion condition 123
- generic options 59
  - router 207–220
  - transport 253–259
- gid (group id)
  - caller 130
  - Exim’s own 171
  - in *queryprogram* router 236
  - in spool file 475
  - local delivery 212
  - of originating user 137
  - system filter 199, 412
- giving up on messages 42
- GnuTLS 338
  - building Exim with 19
  - specifying parameters for 341
- groups
  - additional 213, 255
  - name format 57
  - trusted 204
- gsasl* authenticator 332
- gt** expansion condition 123
- gti** expansion condition 123

## H

- hash function
  - numeric 110, 118
  - textual 106, 117
- header lines
  - added; visibility of 366
  - adding 212
  - adding in an ACL 366
  - adding in transport 254
  - adding; in router or transport 421
  - adding; in system filter 414
  - character sets 106

- header lines (*continued*)
  - Date: 419
  - decoding 106
  - Envelope-to: 419
  - From: 28, 419
  - in expansion strings 106
  - listing 43
  - maximum size of 173
  - Message-ID: 420
  - position of added lines 366
  - position of removed lines 367
  - Received: 420
  - References: 420
  - removed; visibility of 368
  - removing 212, 254
  - removing in an ACL 367
  - removing; in router or transport 421
  - removing; in system filter 414
  - Resent- 418
  - Return-path: 420
  - rewriting 224
  - rewriting at transport time 305
  - Sender: 28, 420
  - transporting 254
  - verifying header names only ASCII 371
  - verifying syntax 371
  - verifying the sender in 371
- header section
  - definition of 5
  - maximum size of 173
- heimdal\_gssapi* authenticator 334
- HELO 424, 446
  - accepting junk data 173
  - ACL for 160, 349, 350
  - argument, setting 295
  - forcing reverse lookup 174
  - forcing use of 296
  - invalid data 427
  - underscores in 173
  - verifying 371
- HELO verifying
  - mandatory 174
  - optional 174
- hex2b64** expansion item 117
- hexquote** expansion item 117
- hiding configuration option values 33, 55
- hints database
  - access by remote transport 259
  - callout cache 386
  - data expiry 190, 315
  - DBM files used for 17
  - deferred deliveries 15
  - ETRN serialization 429
  - maintenance 460
  - not overridden by **-Mc** 42
  - overriding retry hints 41
  - remembering routing 47, 291
  - retry keys 217, 256, 424
  - serializing deliveries to a host 300
  - transport concurrency control 255
  - use for retrying 314
- hmac** 107
- HOME 284
- home directory
  - for local transport 219, 252

- home directory (*continued*)
  - for router 217
- HOST 284
- host
  - ACL checking 369
  - alias for 96
  - error 425
  - for RFC 1413 calls 190
  - limiting SMTP connections from 192
  - list of; randomized 230, 297
  - locally unique number for 179
  - lookup failures 96
  - lookup failures, permanent 97
  - lookup failures, temporary 98
  - maximum number to try 296, 302
  - name in SMTP responses 193
  - name of local 185
  - not logging connections from 175
  - rejecting connections from 175
  - reserved 193
  - serializing connections 300
  - treated as local 176
  - unqualified addresses from 189, 191
  - verifying reverse lookup 372
- host list
  - empty string in 94
  - lookup of IP address 95
  - masked IP address 95
  - matching host name 96, 98
  - matching IP addresses 94
  - mixing names and addresses in 97
  - patterns in 94
  - regular expression in 96
- host name
  - lookup, failure of 133
  - lookup, forcing 175
  - matched in domain list 92
- hosts** ACL condition 369
- HP-UX 172

**I**

- i18n 487
- iconv()* support 19
- id of message 9
- ident *see RFC 1413*
- if**, expansion item 108
- ignoring faulty addresses 249
- imapfolder** expansion item 108
- included address list 241
- inclusions in configuration file 53
- incoming SMTP over TCP/IP 426
- incorporated code 6
- incorrect utf-8 120
- inetd 36, 37, 191, 426
  - wait mode 37
- installing Exim 24
  - info* documentation 25
  - install script options 25
  - testing the script 25
  - what is not installed 24
- integer configuration values 56
- integer format 56
- InterBase
  - server list 176
- InterBase lookup type 75, 86

- interface
  - address, specifying for local message 45
  - listening 149
  - logging 451
  - network 149
- intermittently connected hosts 438
- internationalisation 120
  - email address 487
- IP address
  - binding 298
  - canonical form 117
  - discarding 213
  - for listening 149
  - masking 95, 118
  - normalisation 117
  - testing string format 124
  - translating 218
- IP source routing 473
- ipliteral* router 226
- iplookup* router 227
- iplsearch lookup type 74
- IPv6
  - address scopes 151
  - addresses in lists 57
  - disabling 151, 168
  - DNS black lists 378
  - DNS lookup for AAAA records 169
  - including support for 20
- ip6denorm** expansion item 117
- ip6norm** expansion item 117
- IRIX, DBM library for 17
- isip** expansion condition 124
- isip4** expansion condition 124
- isip6** expansion condition 124

## J

- journal file 11

## K

- keepalive
  - on incoming connection 191
  - on outgoing connection 298
- Kerberos 329

## L

- lc** expansion item 117
- LDAP
  - , 177
  - authentication 83
  - connections 82
  - default servers 178
  - including support for 22
  - lookup type 75
  - lookup, more about 81
  - policy for LDAP server TLS cert presentation 178
  - protocol version, forcing 178
  - query format 81
  - quoting 82
  - returned data formats 85
  - timeout 84
  - TLS cipher suite 177
  - TLS client certificate file 177
  - TLS client key file 177

- LDAP (*continued*)
  - use for authentication 124
  - whether or not to negotiate TLS 178
  - with TLS 81
- ldapauth** expansion condition 124
- le** expansion condition 124
- lei** expansion condition 124
- length** expansion item 108, 117
- length of login name 180
- LF character *see linefeed*
- LHLO argument setting 295
- limit
  - bounce message line length 163
  - bounce message size 164
  - hosts; maximum number tried 302
  - incoming SMTP connections 191
  - message size 181
  - message size per transport 255
  - messages per SMTP connection 192
  - non-mail SMTP commands 191
  - number of hosts tried 296
  - number of MX tried 296
  - number of recipients 189
  - on retry interval 190
  - open files for lookups 180
  - rate of message arrival 195
  - retry interval 314
  - size of message header section 173
  - size of one header line 173
  - SMTP connections from one host 192
  - SMTP syntax and protocol errors 195
  - transport parallelism 255
  - unknown SMTP commands 195
  - user name length 180
- limitations of Exim 3
- limiting client sending rates 378
- line endings 417
- line length
  - maximum 135
- linear search 74
- linefeed 271, 288, 417, 424, 427
- Linux, DBM library for 17
- list
  - address list 99
  - caching of named 91
  - count of items 117
  - domain list 92
  - empty item in 58
  - extracting elements by number 108
  - file name in 90
  - host list 94
  - item count 117
  - iterative conditions 123
  - local part list 101
  - named 90
  - named compared with macro 91
  - negation 89
  - reducing to a scalar 112
  - selecting by condition 106
  - sorting 113
  - syntax of in configuration 57
- list separator
  - changing 57
  - newline as 58
- listcount** expansion item 117

## listextract

extract list elements by number 108

## listing

message body 43

message headers 43

message in RFC 2822 format 43

message log 43

messages on the queue 34

## listnamed expansion item 117

lists of domains; hosts; etc. 89–101

## LMTP

ignoring quota errors 280, 299

logging confirmation 452

over a pipe 280

over a socket 280

over TCP/IP 299, 424

processing details 424–430

*lmtp* transport 280

load average 167, 187, 194

re-evaluating per message 187

## local delivery

definition of 5

using an external agent 289

## local host

domains treated as 176

MX pointing to 217, 222

name of 185

sending to 218, 292

local message reception 32

## local part

ACL checking 370

case of 423

checking in router 214

definition of 5

dots in 423

in retry keys 217

list 101

list, in expansion condition 125

prefix 255, 437

starting with ! 99, 100

suffix 255, 437

local SMTP input 35

## local transports

environment for 251–252

uid and gid 212, 213, 220, 251

local user, checking in router 208

**local\_part** expansion item 118

**local\_parts** ACL condition 370

## local\_scan() function

address rewriting; timing of 305

API description 402

available Exim functions 407

available Exim variables 404

building Exim to use 402

configuration options 403

description of 402–411

memory handling 411

timeout 179

when all recipients discarded 353

Local/eximon.conf 464

Local/eximon.conf 19, 24

Local/Makefile 18, 22

lock files 26, 265

locking files 265, 266, 268, 272, 273

locking mailboxes 462

## log

8BITMIME 159, 450

certificate verification 453

connection rejections 450

cycling local files 443, 458

datedstamped files 443

delayed delivery 450

delivery duration 451

delivery line 215, 446

destination 442

distinguished name 343, 344

DNS failure in list 453

dnslist defer 451

dnssec 451

dropped connection 451

ETRN commands 451

Exim arguments 450

extracts; grepping for 457

**fail** command log line 413

file for each message 10

file path for 180

frozen messages; skipped 452

full parentage 450

general description 442–454

header lines for rejection 452

host lookup failure 451

ident timeout 451

incoming interface 451

incoming proxy address 451

incoming remote port 451

local address and port 451

local files; writing to 443

local interface 451

message log; description of 453

message log; disabling 181

message size on delivery 451

millisecond timestamps 451

non-MAIL SMTP sessions 453

outgoing interface 451

outgoing remote port 451

process ids in 442, 451

process log 186

protocol 446, 487

queue run 451

queue time 451

reception line 445

recipients 452

retry defer 452

return path 452

rewriting 450

selectors 159, 180, 449

sender on delivery 452

sender reception 452

sender verify failure 452

size rejection 452

smtp confirmation 452

SMTP connections 452

SMTP protocol error 453

SMTP syntax error 453

SMTP transaction; incomplete 453

subject 453

summary of fields 448

syslog; writing to 444

tail of; in monitor 465

timezone for entries 180

- log (*continued*)
  - TLS cipher 343, 344, 453
  - TLS peer DN 453
  - TLS SNI 453
  - to file 442
  - to syslog 442
  - types of 442
  - unknown SMTP command 453
  - writing from embedded Perl 148
- log\_message** ACL modifier 359
- log\_reject\_target** ACL modifier 360
- logging
  - custom 489
  - slow lookups 191
- logging in ACL
  - immediate 360
  - specifying which log 360
- LOGIN authentication mechanism 324
- logwrite** ACL modifier 360
- lookup
  - \* added to type 76
  - \*@ added to type 76
  - caching 78
  - cdb 73
  - dbm 73
  - dbm – embedded NULs 73
  - dbm – terminating zero 73
  - dbmjbz 73
  - dbmnz 73
  - default values 76
  - description of 72
  - DNS 75
  - dnsdb 79
  - dsearch 73
  - in domain list 93
  - in expanded string 109
  - inclusion in binary 22
  - InterBase 75, 86
  - iplsearch 74
  - LDAP 75, 81
  - lsearch 74
  - lsearch – colons in keys 74
  - maximum open files 180
  - MySQL 75, 86
  - NIS 74
  - NIS+ 75, 85
  - nwildlsearch 74
  - Oracle 75, 86
  - partial matching 77
  - partial matching – changing prefix 77
  - passwd 76
  - PostgreSQL 76, 86
  - query-style types 75
  - quoting 78
  - Redis 76, 86, 88
  - single-key types 73
  - SQLite 88
  - sqlite 76
  - temporary error in 76
  - types of 73
  - whoson 76
  - wildcard 76, 77
  - wildlsearch 74
- lookup modules 21

- loop
  - caused by **fail** 413
  - in lookups 101
  - prevention 189
  - while file testing 273
  - while routing 14
  - while routing, avoidance of 239
- lower casing 117, 459
- lsearch lookup type 74
- lt** expansion condition 124
- lti** expansion condition 124

## M

- macro
  - compared with named list 91
  - description of 53
  - setting on command line 38
- MAIL
  - ACL for 160, 349
  - logging session without 453
  - rewriting argument of 307
  - SIZE option 427
- mail hub example 234
- mail loop prevention 189
- mailbox
  - locking, blocking and non-blocking 266
  - maintenance 462
  - MMDF format 264
  - multiple 213, 437
  - size warning 271
  - specifying size of 267
  - symbolic link 263, 273
  - time of last read 313
- maildir format 274
  - description of 274
  - maildirfolder* file 268
  - maildirsize* file 267, 276
  - quota; directories included in 267
  - specifying 267
  - time of last read 313
- maildir++ 275
- maildirfolder*, creating 268
- mailing lists 433
  - closed 434
  - for Exim users 2
  - one-time expansion 247
  - re-expansion of 434
  - syntax errors in 433
- mailq* 28
- mailstore format 274
  - description of 276
  - specifying 268
- main configuration 153–206
- main log 442
- maintaining Exim's hints database 460
- malware** ACL condition 370
- malware scan test 33
- manualroute* router 229–235
- mask** expansion item 118
- masked IP address 118
- match** expansion condition 124
- match\_address** expansion condition 125
- match\_domain** expansion condition 125
- match\_ip** expansion condition 125
- match\_local\_part** expansion condition 125

- maximum *see also* limit
- line length 135
- MBX format, specifying 268
- md5** expansion item 118
- MD5 hash 118, 122
- memory
  - debugging 166
- message
  - abandoning delivery attempts 42
  - adding recipients 41
  - age of 135
  - changing sender 42
  - controlling incoming 349
  - copying every 438
  - discarded 447
  - error 425
  - forced failure 413
  - format 32
  - freezing 413
  - frozen 10
  - general processing 416–423
  - header, definition of 5
  - life of 10
  - listing body of 43
  - listing header lines 43
  - listing in RFC 2822 format 43
  - listing message log 43
  - log file for 10, 453
  - manually discarding 42
  - manually freezing 42
  - queueing by file existence 187
  - queueing by load 187
  - queueing by message count 192
  - queueing by SMTP connection count 192
  - queueing certain domains 186
  - queueing remote deliveries 188
  - queueing unconditionally 186
  - reception 9
  - size 136
  - size in queue listing 34
  - size issue for transport filter 300
  - size limit 181
  - submission 365, 416, 419, 420
  - submission, ports for 61
  - thawing frozen 10, 43
  - transporting body only 253
  - transporting headers only 254
- message** ACL modifier 360
  - with **accept** 355
- message body
  - binary zero count 130
  - in expansion 135, 136
  - line count 130
  - newlines in variables 180
  - size 136
  - visible size 181
- message ids
  - details of format 9
  - with multiple hosts 179
- message logs
  - disabling 181
  - preserving 185
- message reference
  - message reference, specifying for local message 46
- message sender, constructed by Exim 9

- Message-ID*: header line 181, 420
- Microsoft Secure Password Authentication 6, 335
- millisecond
  - logging 451
- MIME content scanning 398
  - ACL for 160, 349
  - returned variables 399
- mime\_regex** ACL condition 370
- mixed-case login names 423
- MMDF format mailbox 264
- modifiers
  - dnsdb 79
- monitor *see* Exim monitor
- msglog* directory 453
- multiline responses, suppressing 365
- multiple mailboxes 213, 437
- multiple SMTP deliveries 41, 43, 47, 292, 296, 439
- multiple SMTP deliveries with TLS 346
- multiple spool directories 197
- MX record
  - checking for secondary 223
  - in *dnsdb* lookup 79, 80
  - maximum tried 296
  - not found 223
  - pointing to IP address 162
  - pointing to local host 217, 222
  - required to exist 223
  - security 210, 294
- MySQL
  - lookup type 75, 86
  - server list 182
- N**
- Nagle algorithm 200
- name
  - of local host 185
  - of router 140
  - of sender 40
  - of transport 146
- name server for enclosing domain 80
- named lists 90
- named pipe (fifo) 263
- named queues 48, 138, 188
- ndbm* DBM library 17
- negation
  - in expansion condition 121
  - in lists 89
- network interface 149
- new drivers, adding 491
- newaliases* 28
- newline
  - as list separator 58
  - in message body variables 180
- NFS 216
  - checking for file existence 244
  - lock file 265, 272
- NIS lookup type 74
  - including support for 22
- NIS, retrying user lookups 172
- NIS+ lookup type 75, 85
  - including support for 22
- no\_xxx* *see* entry for *xxx*
- non-immediate delivery 44
- non-SMTP messages
  - ACLs for 159, 349, 350

- not-QUIT, ACL for 161
- NTLM authentication 335
- NUL *see binary zero*
- number of deliveries 190
- numeric comparison 121
- numerical variables (\$1 \$2 etc)
  - in *manualroute* router 231
- numerical variables (\$1 \$2 etc) 128
  - in *cram\_md5* authenticator 327
  - in **errors\_copy** 171
  - in “From ” line handling 418
  - in **gecos\_name** 173
  - in **if** expansion 124
  - in lookup expansion 109
  - in *plaintext* authenticator 323
  - in rewriting rules 306
  - in *spa* authenticator 335
- nwildlsearch lookup type 74

## O

- OCSP
  - stapling 343
- one-time aliasing/forwarding expansion 247
- open files, too many 180
- OpenDMARC 6
- opendmarc
  - acknowledgment 7
- OpenSSL 338
  - building Exim with 19
- OpenSSL
  - compatibility 182
- options
  - appendfile* transport 263
  - authenticator – extracting 34
  - autoreply* transport 277
  - command line 28–50
  - command line; terminating 29
  - configuration – extracting 33
  - cram\_md5* authenticator (client) 327
  - cram\_md5* authenticator (server) 327
  - dnsdb 79
  - dnslookup* router 223
  - generic – definition of 59
  - generic; for authenticators 318
  - generic; for routers 207–220
  - generic; for transports 253–259
  - hiding name of 33
  - hiding value of 33, 55
  - iplookup* router 227
  - lmtp* transport 280
  - macro – extracting 34
  - manualroute* router 229
  - pipe* transport 284
  - plaintext* authenticator (client) 325
  - plaintext* authenticator (server) 323
  - queryprogram* router 236
  - redirect* router 242
  - router – extracting 34
  - smtp* transport 291
  - spa* authenticator (client) 335
  - spa* authenticator (server) 335
  - tls* authenticator (server) 337
  - transport – extracting 34
  - “or” expansion condition 127

- Oracle
  - lookup type 75, 86
  - server list 184
- os.h* 23
- outgoing LMTP over TCP/IP 424
- outgoing SMTP over TCP/IP 424
- ownership
  - alias file 247
  - configuration file 51
  - forward file 247

## P

- packet radio 218
- PAM authentication 126
- pam** expansion condition 126
- panic log 442
- partial matching 77
- passwd file *see /etc/passwd*
- passwd lookup type 76
- PCRE 6, 71
  - security 472
- PCRE library 17
- “percent hack” 184, 389
- periodic queue running 48
- Perl 184
  - calling from Exim 147–148
  - including support for 23
  - standard output and error 148
  - starting the interpreter 47
  - taintmode 147
  - use in expanded string 110
  - use of **warn** 148
- pid (process id)
  - in log lines 442, 451
  - of current process 137
  - of daemon 29, 33, 46
  - re-use of 9
- pid file, path for 185
- pipe
  - duplicated 242
  - in redirection list 240
  - named (fifo) 263
- pipe* transport 282–290
  - , 285
  - environment for command 284, 285
  - failure of exec 285
  - for system filter 199
  - logging output 286
  - path for command 283
  - permitted commands 284
  - returned data 282
  - temporary failure 288
  - uid for 252
  - with multiple addresses 260
- PIPELINING
  - avoiding the use of 296
  - expected errors 195
  - suppressing advertising 185, 365
- pkg-config
  - authenticators 22
  - GnuTLS 20
  - lookups 22
  - OpenSSL 20
- PLAIN authentication mechanism 323
- plaintext* authenticator 323–326

- policy control
  - access control lists 349
  - address verification 382
  - by local scan function 402
  - checking access 459
  - overview 8
  - rejection, returning details 196
  - relay control 389
  - testing 31
- port
  - 465 and 587 61
  - for daemon 166
  - for message submission 61
  - iplookup* router 227
  - logging outgoing remote 451
  - logging remote 451
  - receiving TCP/IP 46
  - sending TCP/IP 299
- PostgreSQL lookup type 76, 86
  - server list 184
- PRDR
  - ACL for 160, 349, 351
  - enabling on server 185
  - enabling, optional in client 298
  - use for per-user SpamAssassin profiles 397
  - variable for 138, 351
- preconditions
  - checking 12
  - definition of 11
  - order of processing 13
- prefix
  - for local part, including in envelope 255
  - for local part, used in router 213
  - for partial matching 77
- preserving domain in redirection 248
- primary host name 92
- printing characters 185
- private options 59
- privilege, running without 471
- privileged user 473
- process id *see pid*
- process log path 186
- process, querying 455
- procmail* 289
- protocol, specifying for local message 46
- proxy
  - access via 485
  - authentication 486
  - client side 485
  - inbound 485
  - logging proxy address 451
  - outbound 485
  - Proxy protocol 485
  - proxy protocol 176
  - server side 485
  - SOCKS 300, 485
- Proxy protocol
  - proxy 485
- proxy support 485
- prvs** expansion item 110
- prvsccheck** expansion item 110
- PTR record
  - in *dnsdb* lookup 79
- public key for signed distribution 3
- pwauthd* daemon 6

- pwcheck* daemon 6, 126
- pwcheck** expansion condition 126

## Q

- query-style lookup
  - definition of 73
  - list of types 75
- querying exim information 32
- queryprogram* router 236–237
- queue
  - count of messages on 34
  - definition of 5
  - delay warning 166
  - delivering specific messages 48
  - display in monitor 466
  - double scanning 47
  - forcing delivery 48
  - grepping 455
  - initial delivery 48
  - listing messages on 34
  - local deliveries only 48
  - menu in monitor 466
  - named 48
  - routing 47
  - summary 456
- queue runner 14, 15, 28, 29
  - abandoning 167
  - definition of 5
  - description of operation 47
  - detecting when delivering from 127
  - for specific recipients 49
  - for specific senders 49
  - logging 451
  - maximum number of 188
  - processing messages in order 187
  - starting manually 47
  - starting periodically 48
- queue\_running** expansion condition 127
- queueing incoming messages 44, 186, 187, 188, 192, 365
- queues
  - named 138, 188
- QUIT, ACL for 161, 349, 351
- quota
  - error testing in retry rule 313
  - imposed by Exim 269
  - in maildir delivery 275, 276
  - maildir; directories included in 267
  - system 262
  - warning threshold 271
- quote** expansion item 118
- quote\_local\_part** expansion item 118
- quoting
  - hex-encoded unprintable characters 117
  - in lookups 78
  - in pipe command 283
  - in regular expressions 119
  - in string expansions 118
  - lookup-specific 118

## R

- Radius 127
- radius** expansion condition 127
- random number 119



- randomized host list 230, 297
- rate limiting 370
  - client sending 378
  - counting unique events 381
  - per\_\* options 379
  - reading data without updating 380
  - strict and leaky modes 380
- RBL *see* DNS list
- RCPT
  - ACL for 63, 161, 349
  - maximum number of incoming 189
  - maximum number of outgoing 299
  - rate limiting 195
  - rewriting argument of 307
  - value of *\$message\_size* 136
- readfile** expansion item 111
- readsocket** expansion item 111
- Received:* header line 420
  - counting 189
  - customizing 188
- receiving mail 9
- recipient
  - ACL checking 370
  - adding 41
  - adding in local scan 405
  - error 426
  - extracting from header lines 49
  - maximum number 189
  - removing 42
  - removing in local scan 405
  - verifying 372
- recipients** ACL condition 370
- redirect* router 238–250
- redirection *see* address redirection
- Redis lookup type 76, 86
- redis lookup type 88
- References:* header line 420
- regex** ACL condition 370
- regex submatch variables (*\$1regex \$2regex* etc) 140
- regular expressions
  - content scanning with 401
  - in address list 99
  - in domain list 93
  - in host list 96
  - in retry rules 311
  - library 6, 71
  - match in expanded string 124
  - quoting 119
  - security 472
- reject log 442
  - disabling 206
- rejection, fake 364
- relaying
  - checking control of 390
  - control by ACL 389
  - testing configuration 31
- remote delivery, definition of 5
- remove\_header** ACL modifier 367
- removing messages 42
- removing recipients 42
- repeated redirection expansion 242
- replacing another MTA 26
- reporting bugs 3
- require** ACL verb 356
- Resent-** header lines 418
  - with **-t** 50
- resolver, debugging output 39
- restricting access to features 166, 186, 245, 246
- retry
  - after long-term failure 315
  - algorithms 314
  - configuration testing 35
  - configuration, description of 310–316
  - default rule 69
  - description of mechanism 15
  - fixed intervals 314
  - increasing intervals 314
  - intermittently working deliveries 316
  - interval, maximum 314
  - limit on interval 190
  - parameters in rules 314
  - quota error testing 313
  - random intervals 314
  - rules 310
  - rules; sender-specific 313
  - specific errors; specifying 312
  - time not reached 310, 449
  - timeout of data 315
  - times 460
- return code
  - for **-bm** 32
  - for **-bS** 35
  - for **-bt** 36
  - for **-bv** 37
  - for **-oee** 44
  - for **-oem** 44
  - for **-oep** 45
  - for bad configuration 51
  - from **run** expansion 112, 140
- return path *see also* envelope sender
  - changing in transport 256
  - created from *Sender:* 421
  - definition of 5
  - in submission mode 417
- Return-path:* header line 256, 420
  - removing 190
- reverse DNS lookup 142, 174, 477
- revocation list 343
- rewriting
  - addresses 10, 304–309, 423
  - at transport time 254, 305
  - bang paths 309
  - flags 307
  - header lines 224
  - logging 450
  - patterns 306
  - replacements 307
  - rules 305
  - testing 35, 305
  - timing of 304
  - whole addresses 308
- RFC 1413 31, 190
  - logging timeout 451
  - query timeout 191
- RFC 2047 19, 308, 400, 409, 422
  - binary zero in 410
  - decoding 119
  - disabling length check 165
  - expansion operator 119

- RFC 3030
- CHUNKING 166, 298, 351
- RFC 7413
- TCP Fast Open 298
- rfc2047** expansion item 119
- rfc2047d** expansion item 119
- rmail* 28
- root privilege 469
- running without 471
- router
  - adding header lines 212
  - carrying on after success 219
  - case of local parts 208
  - changing address for errors 210
  - checking for local user 208
  - checking senders 218
  - customized precondition 208
  - customizing “cannot route” message 207
  - data attached to address 207
  - definition of 11
  - discarding IP addresses 213
  - fallback hosts 211
  - for verification 12
  - forcing verification failure 211
  - go to after “pass” 215
  - home directory for 217
  - IP address translation 218
  - name 140
  - preconditions, order of processing 13
  - prefix for local part 213
  - removing header lines 212
  - requiring file existence 216
  - restricting to specific domains 210
  - restricting to specific local parts 214
  - result of running 12
  - running details 12
  - setting group 212
  - skipping for EXPN 211
  - skipping when address testing 207
  - start at after redirection 215
  - suffix for local part 214
  - timeout 215
  - used only when verifying 220
  - user for filter processing 220
- routers
  - accept* 221
  - default 66
  - dnslookup* 222–225
  - ipliteral* 226
  - iplookup* 227
  - manualroute* 229–235
  - queryprogram* 236–237
  - redirect* 238–250
- routing
  - by external program 236
  - loops in 14, 239
  - whole queue before delivery 47
- rsmtpt* 28
- Rspamd 395
- run** expansion item 112
- run time configuration 51
- runq* 28
- rxquote** expansion item 119

## S

- Samba project 6
- saslauthd* daemon 127
- saslauthd** expansion condition 127
- sasldb2 73
- scanning *see content scanning*
- security
  - build-time features 469
  - command injection attacks 472
  - data sources 472
  - discussion of 469–474
  - local commands 472
  - MX lookup 210, 294
  - regular expressions 472
- sender
  - ACL checking 370
  - address 40, 418
  - authenticated 129
  - changing 42
  - constructed by Exim 9
  - definition of 5
  - gid 137
  - host address, specifying for local message 45
  - host name, specifying for local message 46
  - ident string, specifying for local message 46
  - name 40
  - setting by untrusted user 204
  - source of 36
  - uid 137
  - verifying 372
  - verifying in header 371
- sender\_domains** ACL condition 370
- sender\_retain** submission option 416
- Sender:* header line 28, 420
  - disabling addition of 178
  - retaining from local submission 179
- senders** ACL condition 370
- Sendmail compatibility
  - bi** option 32
  - h** option ignored 40
  - oA** option 43
  - om** option ignored 46
  - oo** option ignored 46
  - t** option 49, 172
  - U** option ignored 50
  - 8-bit characters 29
  - calling Exim as *newaliases* 28
  - command line interface 4
  - “From” line 32, 418, 419
- serializing connections 300
- set** ACL modifier 361
- setuid 469
  - installing Exim with 24
- sg** expansion item 113
- SHA-1 hash 119, 122
- SHA-256 hash 119
- sha1** expansion item 119
- sha256** expansion item 119
- sha3** expansion item 120
- SHA3 hash 120
- shadow transport 256
- Sieve filter 8
  - capabilities 32
  - configuring *appendfile* 262
  - enabling in default router 67

- Sieve filter (*continued*)
  - enabling use of 243
  - forbidding delivery to a file 245
  - “keep” facility; disabling 245
  - not available for system filter 14, 199
  - relative mailbox path handling 262
  - specifying in redirection data 239
  - syntax errors in 249
  - testing 30
  - vacation directory 249
  - value of `$address_file` 129
- SIGHUP 30, 471
- signal exit 285
- SIGUSR1 455
- simultaneous deliveries 251
- single-key lookup
  - definition of 73
  - list of types 73
- size
  - of bounce lines, limit 163
  - of bounce, limit 164
  - of mailbox 267, 271
  - of message 34, 136, 300, 446, 451
  - of message, limit 181, 255
  - of monitor window 465
- SIZE option on MAIL command 424, 427
- skipping faulty addresses 249
- smart host
  - example router 233, 433
  - suppressing queueing 440
- SMTP
  - authentication configuration 317–322
  - batched incoming 35, 430
  - batched outgoing 429
  - batched outgoing; example 234
  - batching over TCP/IP 425
  - callout verification 383
  - command history 143
  - command, argument for 143
  - connection backlog 194
  - connection, ACL for 349, 350
  - delaying delivery 44
  - details policy failures 196
  - encrypted connection 200
  - encryption 338–348
  - error codes 241, 246, 361
  - errors in outgoing 425
  - host name in responses 193
  - incoming call count 193
  - incoming connection count 191, 192
  - incoming over TCP/IP 426
  - input timeout 46, 196
  - limiting incoming message count 192
  - limiting non-mail commands 191
  - limiting syntax and protocol errors 195
  - limiting unknown commands 195
  - listener 29
  - local incoming 429
  - local input 35
  - logging confirmation 452
  - logging connections 452
  - logging incomplete transactions 453
  - logging protocol error 453
  - logging syntax errors 453
  - multiple deliveries 41, 43, 47, 292, 439

- SMTP (*continued*)
  - non-mail commands 428
  - outgoing over TCP/IP 424
  - output flushing, disabling for callout 364, 383
  - output flushing, disabling for delay 364
  - passed connection 41, 43, 47, 292, 425, 439
  - processing details 424–430
  - protocol errors 428
  - rate limiting 195
  - rewriting malformed addresses 307
  - SIZE 257, 300
  - smtps protocol 61, 150, 338
  - smtp protocol 61, 150, 338
  - synchronization checking 194, 364
  - syntax errors 428
  - syntax errors; logging 453
  - testing incoming 31
  - unknown command; logging 453
  - unrecognized commands 428
  - welcome banner 193
- smtp* transport 291–303
- SMTPS 202
- smtps protocol 61, 150, 338
- SMTPUTF8
  - advertising 196
  - ESMTP option 487
  - logging 487
- SOA record
  - in *dnsdb* lookup 79
- socket, use of in expansion 111
- SOCKS
  - proxy 485
- Solaris
  - DBM library for 17
  - flock()* support 272
  - LDAP 81
  - mail* command 40
  - PAM support 126
  - stopping Exim on 27
- sorting
  - a list 113
- sorting remote deliveries 190
- source routing
  - in email address 184
  - in IP packets 473
- SPA authentication 6
- spa* authenticator 335–336
- spam** ACL condition 370
- spam scanning 395
  - returned variables 397
- SpamAssassin 395
- SPF record
  - in *dnsdb* lookup 79
- spool directory
  - checking space 165, 193
  - creating 25
  - definition of 5
  - file formats 197
  - file locked 449
  - files 473
  - files that hold a message 10
  - format of files 475–479
  - input* sub-directory 10
  - path to 197
  - split 197

- “spool file is locked” 452
- spool files
  - editing 475
  - format of 475–479
- sprintf()* 474
- SQL lookup types 86
- sqlite
  - lookup timeout 88
- sqlite lookup type 76, 88
  - lock timeout 198
- src/EDITME* 18
- SRV record
  - enabling use of 223
  - in *dnsdb* lookup 79
  - required to exist 223
- SSL *see* TLS
- SSMTP 202
- smtp protocol 61, 150, 338
  - outbound 299
- STARTTLS, ACL for 161, 349
- stat** expansion item 120
- statistics 458
- statvfs** function 465
- “sticky” bit 26, 265
- str2b64** expansion item 120
- string
  - case forcing 117, 120
  - comparison 123, 124
  - expansion *see* expansion
  - format of configuration values 56
  - length in expansion 120
  - list, definition of 57
  - quoted 56
  - testing for IP address 124
- stripchart 464
- strlen** expansion item 120
- subject, logging 453
- submission fixups, suppressing 365
- submission fixups, suppressing (command-line) 40
- submission mode 365, 416, 419, 420
- substr** expansion item 113, 120
- substring extraction 113, 120
- suffix for local part
  - including in envelope 255
  - used in router 214
- SUPPORT\_TLS 19
- symbolic link
  - to build directory 17
  - to *exim* binary 26
  - to mailbox 263, 273
  - to source files 21
- synchronization checking in SMTP 194, 364
- syntax of common options 55
- syslog 442
  - duplicate log lines; suppressing 198
  - facility; setting 198
  - pid 198
  - process name; set with flag 41
  - process name; setting 198
  - timestamps 199
- system aliases file 24
- system filter 412–415
  - specifying 199
  - testing 30
- system log 442

## T

- TCP Fast Open
  - enabling, in client 298
- TCP\_NODELAY on listening sockets 200
- TCP\_WRAPPERS\_DAEMON\_NAME 20
- tcp\_wrappers\_daemon\_name* 20
- TCP/IP
  - logging incoming remote port 451
  - logging local address and port 451
  - logging outgoing remote port 451
  - logging proxy address 451
  - setting listening interfaces 46, 149
  - setting listening ports 46, 149, 166
  - setting outgoing port 299
- tcpwrappers, building Exim to support 20
- tdb* DBM library 18
- terminology definitions 4
- testing
  - , 33
  - addresses 36, 211
  - filter file 30
  - forward file 30
  - incoming SMTP 31
  - installation 25
  - relay control 31
  - retry configuration 35
  - rewriting 35, 305
  - string expansion 30, 43, 102
  - system filter 30
  - variables in drivers 209, 253
- thawing messages 43, 162, 467
- time interval
  - decoding 120
  - formatting 120
  - specifying in configuration 56
- time\_eval** expansion item 120
- time\_interval** expansion item 120
- timeout
  - dns lookup 80, 169
  - for *local\_scan()* function 179
  - for non-SMTP input 46, 188
  - for RFC 1413 call 191
  - for SMTP input 46, 196
  - frozen messages 200
  - LDAP lookup 84
  - mailbox locking 266
  - of retry data 315
  - of router 215
  - SQLite 88
- timezone, setting 200
- timestamps
  - millisecond, in logs 451
- TLS
  - advertising 200
  - automatic start 50
  - avoiding for certain hosts 296
  - broken clients 203
  - certificate status 202
  - client certificate revocation list 300
  - client certificate verification 203, 204, 343, 370
  - client certificate, location of 300
  - client private key, location of 301
  - configuring an Exim client 344
  - configuring an Exim server 341
  - D-H bit count 201

- TLS (*continued*)
  - D-H parameters for server 201
  - Diffie-Hellman minimum acceptable size 300
  - EC cryptography 202
  - esmtplib state; remembering 203
  - including support for TLS 19
  - logging cipher 453
  - logging peer DN 453
  - logging SNI 453
  - multiple message deliveries 296, 346
  - OCSP proof file 202
  - on SMTP connection 338
  - OpenSSL vs GnuTLS 338
  - passing connection 296
  - requiring for certain servers 297
  - requiring specific ciphers 203, 301
  - requiring specific ciphers (OpenSSL) 340
  - revoked certificates 343
  - server certificate hostname verification 301
  - server certificate revocation list 201
  - server certificate verification 301, 302
  - server certificate; location of 200
  - Server Name Indication 146, 301, 345
  - server private key; location of 203
  - specifying ciphers (GnuTLS) 341
  - specifying key exchange methods (GnuTLS) 341
  - specifying MAC algorithms (GnuTLS) 341
  - specifying priority string (GnuTLS) 341
  - specifying protocols (GnuTLS) 341
  - SSL-on-connect outbound 299
  - use without STARTTLS 50
- tls* authenticator 337
- tmllib** 286
- To:* header line 49
- too many open files 180
- top bit *see* 8-bit characters
- tr** expansion item 114
- trailing dot on domain 198
- transport
  - body only 253
  - concurrency limit 255
  - current directory for 253
  - definition of 11
  - external 4
  - filter 132, 137, 257, 283, 300, 424
  - filter, timeout 258
  - generic options for 253–259
  - group; additional 255
  - group; specifying 254
  - header lines only 254
  - header lines; adding 254
  - header lines; removing 254
  - header lines; rewriting 254
  - home directory for 255
  - local 212, 213, 220
  - local; address batching in 260
  - local; current directory for 252
  - local; environment for 251–252
  - local; home directory for 252
  - local; uid and gid 251
  - message size; limiting 255
  - name 146
  - parallel processes 255
  - return path; changing 256
  - shadow 256

- transport (*continued*)
  - user, specifying 258
- transports
  - appendfile* 262–276
  - autoreply* 277–279
  - default 68
  - lmtpl* 280
  - pipe* 282–290
  - smtp* 291–303
- Tru64-Unix build-time settings 22
- trusted groups 204
- trusted users 40, 204, 473
  - definition of 28
- TTL
  - of dns lookup 80
- TXT record
  - in *dnsdb* lookup 79
- U**
- uc** expansion item 120
- UDP communications 361
- uid (user id)
  - caller 130
  - Exim's own 171
  - for *queryprogram* 236
  - in spool file 475
  - local delivery 220, 258, 287
  - of originating user 137
  - system filter 199, 412
  - unknown caller 204
- underscore in EHLO/HELO 173
- unfreezing messages 43, 162, 467
- Unicode 117
- unknown host name 96, 98
- unprivileged delivery 167
- unprivileged running 471
- unqualified addresses 189, 191, 417
- untrusted user setting sender 204
- upgrading Exim 27
- upper casing 120
- USE\_DB 18, 460
- USE\_GNUTLS 20
- USE\_TCP\_WRAPPERS 20
- user
  - admin 473
  - admin definition of 28
  - trusted 40, 204, 473
  - trusted definition of 28
  - untrusted setting sender 204
- user name
  - format of 57
  - maximum length 180
- UTF-8
  - conversion from 117
  - expansion 120
  - in domain name 162
  - mail name handling 487
- utf-8
  - utf-8 sequences 120
- utf8\_domain\_from\_alabel** expansion item 120
- utf8\_domain\_to\_alabel** expansion item 120
- utf8\_localpart\_from\_alabel** expansion item 120
- utf8\_localpart\_to\_alabel** expansion item 120
- utf8clean** expansion item 120
- utilities 455–463

## UUCP

example of router for 235  
“From” line 32, 176, 205, 418

## V

vacation processing 438  
Variable Envelope Return Paths 435  
variables *see expansion, variables*  
variables (*\$auth1 \$auth2* etc)  
    in *tls* authenticator 337  
**verify** ACL condition 370, 371, 372  
verifying  
    EHLO 371  
    header names only ASCII 371  
    header syntax 371  
    HELO 371  
    not blind 371  
    recipient 372  
    redirection while 387  
    sender 372  
    sender in header 371  
    suppressing error details 386  
verifying address  
    by callout 383  
    differentiating failures 382  
    options for 382  
    overview 12  
    using **-bv** 36  
VERP 435  
version number of Exim 36  
virtual domains 436  
virus scanners  
    avast 392  
    clamd 392  
    command line interface 393  
    DrWeb 393  
    f-prot6d 393  
    f-protd 393  
    F-Secure 394  
    Kaspersky 392, 394  
    mkscd 394  
    simple socket-connected 394  
    Sophos and Sophie 394  
virus scanning 391  
VRFY  
    ACL for 161, 349  
    error text, display of 241  
    processing 428

## W

**warn** ACL verb 356, 367  
    log when skipping 450  
warning of delay 166  
    customizing the message 205, 432  
web site 2  
welcome banner for SMTP 193  
white space  
    in configuration file 52  
    in header lines 106  
    in *lsearch* key 74  
whoson lookup type 76  
wiki 2  
wildcard lookups 76, 77  
wildsearch lookup type 74

window size 465

## X

*X-Failed-Recipients*: header line 15  
X-windows 7, 464  
X11 libraries, location of 23

## Z

zero, binary *see binary zero*