JUCE

Jam Unsolicited Commercial Email

Please select one of the topics below for configuration information

Introduction

Junk mail is now becoming a serious threat to the commercial use of the Internet and so we have introduced this software to ensure that users of NTMail software do not suffer.

By using this software you should be able to eradicate the presence of Junk Mail from your system almost entirely. It will certainly ensure that no-one can use your system as a relay for their junk mail and thus reserve all your resources and bandwidth for your own use.

There are features included to stop mail containing specific keywords which will stop all mail containing these words from entering your system as well as more server specific facilities. These include the ability to only allow connections from certain IP addresses, deny connections from certain IP addresses, to verify that connecting servers are who they say they are, to stipulate the IP addresses that remote servers must have before they are allowed to connect and a host of other features.

Perhaps the one that is going to be the most successful is the Artificial Intelligence feature as, no matter how well you protect your server, junk mailers are a resourceful lot and they will do their best to get round whatever protection you put in place. The AI feature will maintain an ongoing watch on the traffic passing through your system, spot any traffic that varies from the norm, and take steps to prevent this traffic from entering your system.

To use this software effectively, you will need to be familiar with the way that email is delivered to your server using the protocol called SMTP. If you are not familiar with this protocol, please take time to read the SMTP Primer.

You can navigate through the various screens available to you either by selecting the tabs at the top of the screen or by using the Back and Next buttons displayed on each screen.

Once you have completed entering all the information you want you should click on the Apply button to confirm the changes and then on the OK button to close the program. Clicking on Cancel will close the program without saving any changes.

Artificial Intelligence

The Artificial Intelligence feature will help to catch any unusual activities on your server, normally as the result of an unauthorised person using you as a relay server. It is a failsafe feature that usually requires no configuration but only acts in extreme circumstances. The AI feature acts on each RCPT command of the SMTP protocol and before a message is accepted for delivery by the server.

NTMail will monitor the messages passing through the server and count the number from a given email address, to a given email address and from an IP address. Over a period of time (Required Samples) the server will have built up a profile of the messages that pass through it under normal conditions. Once the profile has been created, the server will check to see that the number of messages for that mail address in any particular day does not exceed the average number of messages per day multiplied by the "Average Multiplier". The "Running Average Minimum" defines the minimum that the running average can be for any address.

AI Enabled for MAIL Clause

Required Samples
Enter a figure to indicate the number of days the AI feature should sample your mail server to build up a profile of your mail throughput.

Average Multiplier
This entry allows you to set the multiplier that will be applied to the running average before the AI features start rejecting mail, i.e. if you set a multiplier of 2 and the running average for a particular user is 3, the maximum number of messages allowed will be 6 in any one day. Thus the 7th message will be rejected. As time progresses, the average will adjust at a maximum rate of "average multiplier"/"required samples" messages per day towards a new upper band.

Running Average Minimum
This parameter tells the mail server that any new account will have the maximum threshold specified here. On setting up a new account the AI feature will allow a maximum throughput of "Running Average Minimum" multiplied by "Average Multiplier" messages to that account before kicking in. This is necessary due to the fact that a new account will never have sent or received mail previously. A remote server sending mail exceeding this threshold will have to wait until the next day to deliver any more mail.

Log Entry
A typical log entry might look as follows:

SPAM	24 Sep 97	16:37:41	H	3296	14	AI IP	Response (25 average - now 250) "123.123.123.123
SPAM	24 Sep 97	16:37:41	H	3290	14	Send Response	user@another.co.uk user@mail.net -> "453 Exceede

Failure Message
Any mail that fails to pass the AI checks should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

AI Enabled for RCPT Clause

As for MAIL Clause above.

AI Enabled for IP Address

As for MAIL Clause above.

The AI checks are performed in the first place on the IP address of the remote machine, if this passes they are performed on the MAIL clause and if this passes then on the RCPT clause. If any one of these checks fails the further checks are not carried out. The fields are initialised with the following values which we would recommend as an optimum setting although you may wish to change them to better reflect your own setup.

Parameter	**Value**
AIMAILRunningAverageMin	20
AIMAILAverageMultiplier	10
AIRCPTRunningAverageMin	20
AIRCPTAverageMultiplier	10
AIIPRunningAverageMin	20
AIIPAverageMultiplier	8

DATA

The DATA page allows you to define which domains are allowed to be served through your server. If you are acting as a backup or relay for another server please be sure to enter them in the **Treat as Local** section. This feature takes effect after a remote server has connected to you and informed you of both who the mail is from and who it is addressed to but before the message itself is transferred.

Only Accept Local

Selecting the Only Accept Local check box will deny access to your server to all mail except that either addressed to or being sent from one of your local domains. This option automatically recognises all local domains including POP domains.

Please Note: This does not prevent a remote system pretending mail was sent from your domain. To do this you will need to use the Local IP option.

Treat as Local

If your server is acting as a backup or relay server for any non-local domains they must be added to the Treat as Local section of this page. To add an entry type the domain name you want to be treated as local in the text box on the left and then click on the right arrow to add it to the list. Similarly to remove a domain, highlight it in the list on the right and click on the left arrow. Wildcards may be used.

Reject with

Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

Log Entry

A typical log entry might look as follows:

```
SPAM    24 Sep 97    16:37:4     H      3298    14    OnlyAcceptLocal        rejected j@nts,.co.uk from [196.198.250.250] [19
                     1
```

End Of Message

The End of Message tab allows you to restrict the maximum incoming and outgoing messages sizes that will be allowed to pass through NTMail on a per domain basis. As you would imagine this feature is only applied to messages after the full message has been transferred to your server and the remote server has sent an "End of Message" signal.

Maximum Outbound Message Size

Domain
Enter the name of the domain that this setting will be applied to.

Size
Enter the value in Kb of the maximum allowed outgoing message size.

Failure Message
Any mail that fails should tell the sender why it is being rejected otherwise they will continue to try to send it to you. Failure messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

Once you have entered the details for a domain, click on the right arrow to move them over to the list on the right. To remove or edit the Maximum Outbound Message Size for a domain, highlight the domain in the list on the right and click the left arrow button.

Maximum Inbound Message Size

As for Maximum Outbound Message size described above except it affects incoming messages.

Log Entry
A typical log entry might look as follows:

SPAM	24 Sep 97	16:37:41	H	3301	14	MaxMessageSize	rejected 311 kBytes to user@mail.net from us [196.198.250.250]
SPAM	24 Sep 97	16:37:41	H	3302	2	MaxMessageSizeOut	rejected 311 kBytes from user@mail.net to us [196.198.250.250]

HELO

Remote IP

The Remote IP option allows you to enter the IP address of machines you are willing to accept connections from, it will then verify that a connecting machine is from one of these IP addresses or refuse the connection.

Enter IP Address

Enter the IP Address in the text box and then click the right arrow to add it to the list on the right. Similarly to remove an IP Address highlight it in the list on the right and click on the left arrow.

Reject with

Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

Machine Name

The Machine Name section similarly performs checks to help ensure that the connecting machine is who it says it is. It will take the IP address of a connecting machine and do a reverse lookup on it. If this does not match the name in the HELO command, the connection is rejected. For example, if the remote machine says "HELO mail.net-shopper.co.uk" and the result of the lookup does not match then the connection will be rejected.

Use Raw IP Address in Logs

Check this option if you would like the raw IP address of the machine used in the logs rather than the machine name.

Reverse Lookup on IP and Terminate if not the Same

Checking this option will cause NTMail to do a reverse DNS Lookup on the IP address of the connecting machine and terminate the connection if the results do not match. For example, if the remote machine says "HELO mail.net-shopper.co.uk" and the IP address 194.205.1.152 resolves to mail.net-shopper.co.uk then the connection would be accepted.

Reject with

Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

Use Result of Reverse Lookup in Logs

Checking this option tells NTMail to use the result gathered from the reverse lookup in the logs rather than the information given out by the connecting machine.

Please Note: Some servers may have a "non-existant" reverse lookup, if this is the case the real source of the message will be lost.

Log Entry

A typical log entry might look as follows:

SPAM	24 Sep 97	16:37:41	H	3289	14	HELO	IPAllowed - rejected [196.198.250.250] [196.198.250.250]
SPAM	24 Sep 97	16:37:41	H	3290	2	HELO	VerifyHostname with IP Address- inconsistent relay.mail.ne

MAIL

Local IP

The Local IP option provides the facility to specify the IP address that sending mail servers for your local domains must use before they are allowed to post mail with a local MAIL clause through your server. This facility is particularly useful in preventing attacking systems pretending to be local so as to bypass the Only Accept Local security feature. The check is done on the MAIL clause which immediately follows the HELO clause and so these messages never actually reach your server.

Sending Domain
Select a local domain that this setting is to be applied to from the drop down list.

IP Address
Specify a list of allowed IP Addresses for this domain by entering them one at a time in the text area to the left and clicking on the right arrow to add them to the list. They can be removed by highlighting the desired IP Address in the list on the right and clicking on the left arrow.

Reject with
Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

Reverse MX Check

This option provides an added security feature that allows for verification of a sending server.

Do Reverse Lookup
Selecting this option will cause the SMTP server to do a lookup on the domain specified in the MAIL clause, if this does not resolve then the mail is rejected.

Reject with
Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

Log Entry
A typical log entry might look as follows:

SPAM	24 Sep 97	16:37:41	H	3293	24	MAIL	AllowedSenderIP - rejected johns@net-shopper.co.u
SPAM	24 Sep 97	16:37:41	H	3294	25	VerifyIncoming	rejected user@mail.net from [196.198.250.250] [196

RCPT

RCPT Clauses

This page allows you to specify how mail with multiple recipients is handled. It is an ideal method of dealing with Spam generated by sending a single mail message to many people.

Reasons this value should be high:

- Many legitimate mail servers will use the multiple RCPT clause to reduce network traffic for a common message - for example a list server.
- Mail clients often use the multiple RCPT clause to send "CC" and "BCC" copies of messages.
- Some messaging systems are incapable of recovery if a failure occurs part way through a list of RCPT clauses.

Reasons why this value should be low:

- Someone wishing to hijack your mail server would like to deliver a message to you once and have you deliver thousands of copies on their behalf.

We recommend a compromise value of 5 or 10 for the maximum number of RCPT clauses.

Default Maximum RCPT Clauses
This sets the default maximum RCPT clauses allowed for a domain. RFC 821 specifies that this should be no greater than 100 but you can set it to whatever you want.

Domain
Select the domain that you want all the other settings on this page to be applied to from the drop down list.

IP
Mail from specific domains may be set to use RCPT clause limits other than the default. For instance, you may want to specify that any mail from a server with an IP Address of 196.198.12.12 can have a maximum of 5 RCPT clauses. In this case you would enter 196.198.12.12 in this box and the number of recipients in the RCPT Clauses box and then click the right arrow to add the entries to the list on the right. Similarly to remove an entry highlight it in the list area and click the left arrow. For example if there were 3 entries each with a maximum RCPT limit of 5 and one with 20 the entries would look like this:

196.198.12.12:5
196.198.12.13:5
196.198.12.14:5
196.198.12.15:20

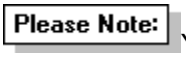

You must use the IP Address and not the name of the server.

RCPT Clauses
Specifies the maximum number of RCPT clauses allowed per mail message to be applied to the domain entered above.

Reject with
Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

Reverse MX Check

Do MX Lookup on each address
Selecting this option will cause the SMTP server to do a lookup on the domain specified in the RCPT clause, if this does not resolve then the mail is rejected.

Reject with
Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

Log Entry
A typical log entry might look as follows:

SPAM	24 Sep 97	16:37:41	H	3293	10	VerifyIncoming	rejected user@mail.net from [196.198.250.250] [196.198.250.250]
SPAM	24 Sep 97	16:37:41	H	3140	10	RcptLimit	user@domain.com from relay.mail.net [8720156]

Redirect

Example

Mail to or from specific locations may be redirected depending on the entries made on this page. Wildcards can be used to include complete domains. The entire redirect settings are parsed for each message so an earlier setting may be overridden by one further down in the list.

From

Enter the wildcarded mail address or IP address that incoming mail from will be be used to initiate redirection of messages, i.e. *@spammer.com would affect all messages that come from the domain spammer.com. If the "Use IP Address" option is selected, an IP address can be used instead of an email address, e.g. 192.192.5.8/8 would affect all messages from the C Class address 192.192.5.0 -> 192.192.5.255

Use IP Address

Checking the **Use IP Address** checkbox at the top of the page will enable this area into which you should enter the IP address of a server you want the redirect function to operate on. Operations are the same as for the MAIL clause which would normally be used.

To

Enter the name that the mail message is addressed to.

Operations

There are three operations that can be carried out on incoming mail, redirect, don't redirect (i.e no action) and respond. These options would normally be carried out on email addresses, however IP Addresses may also be used by checking the **Use IP Address** option.

Redirect

Redirect does exactly as its name implies, it takes the message destined for one account and redirects it to another. This redirection could be to any account. The special account NULL is used when you simply want to throw messages away, i.e. the message is accepted to your server but is then simply deleted.

Don't Redirect

Don't redirect would normally be used to stop the redirection of messages for one or more users in a domain where all the users have previously been redirected. For instance you may want to redirect all mail from a domain to NULL except for that from the PostMaster.

Respond

This option allows you to enter a response message that should be returned to the mail server of anyone posting to/from this particular user or group of users.

Redirect To

Enter an account that you want the mail message redirected to, this should a local address.

Respond With

Enter a response message. For instance, the response could read "525 User does not want to accept mail from this domain", as for rejection messages SMTP Reply Codes should be used in the message.

Add Entry

Once you have completed all the information required above, clicking on the Add Entry button will cause it to be entered into the redirect file.

Remove

The opposite of the Add Entry button, highlighting an action in the list and clicking on Remove will remove

that entry from the redirect file.

Comments
Comments may be attached to particular entries by highlighting the entry in the list, clicking on the Comments button and entering the text you want to use as a comment. This is particularly useful if you wish to remember why you decided to redirect a particular set of email.

Up
Highlighting an entry in the list and clicking the **Up** button will move this entry further up the redirect file.

Down
Highlighting an entry in the list and clicking the **Down** button will move this entry further down the redirect file.

Log Entry
Typical log entries might look as follows:

SPAM	24 Sep 97	16:37:41	H	3140	9	redirect	user@another.co.uk ([196.198.250.250]) usera@ma
SPAM	24 Sep 97	16:37:41	H	H	4	Send Response	user@another.co.uk user@mail.net -> "550 Unacce
SPAM	24 Sep 97	16:37:41	H	3140	3	redirect	user@another.co.uk user@mail.net -> NULL

Restricted Word File

Example

Domain

Select the Domain that you want the restricted words applied to from the drop down list.

Phrase

Enter a word or phrase that you want to disallow from mail messages then click the right arrow to add them to the list. Care should be taken to select the words carefully as all mail containing these words will be rejected. If using a complete word you should follow it with a space otherwise it will restrict words based on partial matches. To remove a word, highlight it in the list and click the left arrow.

Reject with

Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details.

Save Copy to SpamMes Directory

Saves a copy of any rejected mail to the domain\SpamMes directory. It is worthwhile reviewing the messages in this directory from time to time to ensure that genuine mail is not being rejected.

Return Error to Remote System Administrator

Returns an error message to the remote System Administrator telling them that their email has been rejected due to the message contents.

Log Entry

A typical log entry might look as follows:

SPAM 24 Sep 97 16:37:41 H 3302 14 RestrictedWordList rejected user@another.co.uk from [196.198.250.2

SMTP Reply Codes

Theory of Reply Codes

A full explanation of SMTP Reply Codes may be found in RFC 821 - Theory of Reply Codes. Here is a quick summary:

All responses have 3 digits followed by a space and a free-format text response.

- Responses starting "2" indicate success.
- Responses starting "4" indicate a transient failure. Retrying later may succeed.
- Responses starting "5" indicate a permanent failure. The message will never be accepted and should be returned.

Examples of reply codes as used by JUCE in the various stages of the SMTP Protocol:

Clause	Code	Message
HELO	**453**	Exceeded IP count - please try later
	550	Your server has been banned from this server
	450	Too many messages from you today
MAIL	**453**	Exceeded MAIL count - please try later
	550	Domain has no MX or ANAME record
RCPT	**550**	Too many RCPT clauses
	450	Too many messages to this user today
	453	Exceeded RCPT count - please try later
DATA	**550**	This mail is not local
	550	You are not allowed to post to this address
End of	**452**	Insufficient system storage
Message	**552**	Exceeded maximum message size

Theory of Reply Codes

The three digits of the reply each have a special significance. The first digit denotes whether the response is good, bad or incomplete. An unsophisticated sender-SMTP will be able to determine its next action (proceed as planned, redo, retrench, etc.) by simply examining this first digit. A sender-SMTP that wants to know approximately what kind of error occurred (e.g., mail system error, command syntax error) may examine the second digit, reserving the third digit for the finest gradation of information.

There are five values for the first digit of the reply code:

1yz Positive Preliminary reply
The command has been accepted, but the requested action is being held in abeyance, pending confirmation of the information in this reply. The sender-SMTP should send another command specifying whether to continue or abort the action.
[Note: SMTP does not have any commands that allow this type of reply, and so does not have the continue or abort commands.]

2yz Positive Completion reply
The requested action has been successfully completed. A new request may be initiated.

3yz Positive Intermediate reply
The command has been accepted, but the requested action is being held in abeyance, pending receipt of further information. The sender-SMTP should send another command specifying this information. This reply is used in command sequence groups.

4yz Transient Negative Completion reply
The command was not accepted and the requested action did not occur. However, the error condition is temporary and the action may be requested again. The sender should return to the beginning of the command sequence (if any). It is difficult to assign a meaning to "transient" when two different sites (receiver- and sender- SMTPs) must agree on the interpretation. Each reply in this category might have a different time value, but the sender-SMTP is encouraged to try again. A rule of thumb to determine if a reply fits into the 4yz or the 5yz category (see below) is that replies are 4yz if they can be repeated without any change in command form or in properties of the sender or receiver. (E.g., the command is repeated identically and the receiver does not put up a new implementation.)

5yz Permanent Negative Completion reply
The command was not accepted and the requested action did not occur. The sender-SMTP is discouraged from repeating the exact request (in the same sequence). Even some "permanent" error conditions can be corrected, so the human user may want to direct the sender-SMTP to reinitiate the command sequence by direct action at some point in the future (e.g., after the spelling has been changed, or the user has altered the account status).

The second digit encodes responses in specific categories:

x0z Syntax -- These replies refer to syntax errors, syntactically correct commands that don't fit any functional category, and unimplemented or superfluous commands.

x1z Information -- These are replies to requests for information, such as status or help.

x2z Connections -- These are replies referring to the transmission channel.

x3z Unspecified as yet.

x4z Unspecified as yet.

x5z Mail system -- These replies indicate the status of the receiver mail system vis-a-vis the requested transfer or other mail system action.

The third digit

The third digit gives a finer gradation of meaning in each category specified by the second digit. The list of replies illustrates this. Each reply text is recommended rather than mandatory, and may even change according to the command with which it is associated. On the other hand, the reply codes must strictly follow the specifications in this section. Receiver implementations should not invent new codes for slightly different situations from the ones described here, but rather adapt codes already defined.

For example, a command such as NOOP whose successful execution does not offer the sender-SMTP any new information will return a 250 reply. The response is 502 when the command requests an unimplemented non-site-specific action. A refinement of that is the 504 reply for a command that is implemented, but that requests an unimplemented parameter.

The reply text may be longer than a single line; in these cases the complete text must be marked so the sender-SMTP knows when it can stop reading the reply. This requires a special format to indicate a multiple line reply.

Multi-line Replies

The format for multi-line replies requires that every line, except the last, begin with the reply code, followed immediately by a hyphen, "-" (also known as minus), followed by text. The last line will begin with the reply code, followed immediately by , optionally some text, and .

For example:

123-First line
123-Second line
123-234 text beginning with numbers
123 The last line

In many cases the sender-SMTP then simply needs to search for the reply code followed by at the beginning of a line, and ignore all preceding lines. In a few cases, there is important data for the sender in the reply "text". The sender will know these cases from the current context.

The Log File

This page describes how you can read the log file associated with JUCE so you can better see what mail is being rejected by your server and for what reason. Every action taken by JUCE will result in a line being entered in the SL??????.LOG file located in your ntmail\log directory, ?????? indicates the date that the particular log corresponds to.

Each line in the log will be of the form:

Process	Date	Time	Log ID	Error Code	Thread ID	JUCE Action	Mail Actio
SPAM	The Current Date	The Current Time	H	Numeric Code	Numeric Code	The action that caused the entry	What happ to the mail message

The first three of these options are self-explanatory, the others are described below.

Log ID
The Log ID should always be equal to H for this process and is the internal NTMail code that is used to assign the messages to the correct log.

Error Code
This is a 4 digit error code for the use of Internet Shopper Ltd to assist with tracking Observation Reports.

Thread ID
The Thread ID is the thread number of the mail message that caused the entry in the Log File. The message will also appear in the SMTP Log with the same Thread ID.

JUCE Action
The action indicated here is the function that JUCE took on this message, i.e. redirect, Send Response, reject, etc.

Mail Action
The Mail Action gives details of the message itself including who it was from, who it was addressed to and what happened to it.

For example, a typical log entry as a result of an entry in the redirect file might look as follows:

SPAM	24 Sep 97	16:37:41	H	3140	3	redirect	john@net-shopper.co.uk brian@net-shopper.co.uk -> NULL

This shows that as a result of an entry in the redirect file a message from john@net-shopper.co.uk to brian@net-shopper.co.uk was sent to the special account NULL, essentially the mail was deleted.

Full details of each logged response may be found under the relevant section:

- HELO
- MAIL
- RCPT
- Redirect File
- Artificial Intelligence
- DATA
- End Of Message
- Max Messages
- Restricted Word File

RFC 821 - Simple Mail Transfer Protocol

Below are a range of relevant extracts from RFC 821 which will hopefully assist you in deciding how best to use the features available in JUCE.

SMTP Primer

The SMTP Protocol is fairly straightforward although it may not seem so at first. Here we will walk you through the various stages of the protocol.

When a remote server first connects to your SMTP server your's will issue a ready command that looks something like this.
220 mail.yourserver.com WindowsNT SMTP Server v3.03.003 ESMTP ready at Thu, 18 Sep 1997 00:47:02 +0100

The Remote server will then issue a HELO command and the Command sequence will be worked through until the message has been delivered and the connection closed down. A typical scenario is shown below:

220 ntmail.local.com WindowsNT SMTP Server v3.03.0003 ESMTP ready at Tue, 16 Sep 1997 08:58:40 +0100

HELO server.remote.com

250 ntmail.local.com server.remote.com

MAIL FROM:<brian@remote.com>

250 Ok.

RCPT TO:<john@local.com>

250 Ok.

DATA

354 Start mail input, end with <CRLF>.<CRLF>.

... message sent ...

250 Requested mail action Ok.

QUIT

221 Goodbye server.remote.com

As you can see from the above the connecting (remote) server is in charge of the mail transfer but your local server is given the opportunity to reply to each of the commands issued before the transmission progresses. By issuing the correct responses your service can allow the transmission to continue or it can halt it either temporarily or permanently. Each response that your server sends to the remote server must begin with a 3 digit SMTP Reply Code, it is this code that governs what happens next and although not essential it is customary to append some information to each of the Reply Codes used.

In a typical SMTP transaction this gives you no less than 5 opportunities to either accept or reject messages.

IP Addresses

There are many areas of this software where you will need to use IP addresses rather than domain names for entering information. Here, we explain the various ways in which IP addresses may be entered.

a.b.c.d	Specific IP address
a.b.c.*	All IP addresses beginning a.b.c
a.b.c.d-e	A range of IP addresses from d to e
a.b.c.d/n	Ignore n bits

An "!" may be placed at the beginning of the address to indicate NOT.

Examples

!194.194.194.194	NOT IP Address 194.194.194.194
194.194.194.*	Addresses in the range 194.194.194.0 -> 194.194.194.255
194.194.194.194/10	Addresses in the range 194.194.192.0 -> 194.194.195.255
194.194.192-194.*	As above
194.194.194.194/16	Addresses in the B Class range 194.194.0.0 -> 194.194.255.255

RFC 821 Introduction

The objective of Simple Mail Transfer Protocol (SMTP) is to transfer mail reliably and efficiently.

SMTP is independent of the particular transmission subsystem and requires only a reliable ordered data stream channel. Appendices A, B, C, and D describe the use of SMTP with various transport services. A Glossary provides the definitions of terms as used in this document.

An important feature of SMTP is its capability to relay mail across transport service environments. A transport service provides an inter process communication environment (IPCE). An IPCE may cover one network, several networks, or a subset of a network. It is important to realize that transport systems (or IPCEs) are not one-to-one with networks. A process can communicate directly with another process through any mutually known IPCE. Mail is an application or use of inter process communication. Mail can be communicated between processes in different IPCEs by relaying through a process connected to two (or more) IPCEs. More specifically, mail can be relayed between hosts on different transport systems by a host on both transport systems.

RFC 821 The SMTP Model

The SMTP design is based on the following model of communication: as the result of a user mail request, the sender-SMTP establishes a two-way transmission channel to a receiver-SMTP. The receiver-SMTP may be either the ultimate destination or an intermediate. SMTP commands are generated by the sender-SMTP and sent to the receiver-SMTP. SMTP replies are sent from the receiver-SMTP to the sender-SMTP in response to the commands.

Once the transmission channel is established, the SMTP-sender sends a MAIL command indicating the sender of the mail. If the SMTP-receiver can accept mail it responds with an OK reply. The SMTP-sender then sends a RCPT command identifying a recipient of the mail. If the SMTP-receiver can accept mail for that recipient it responds with an OK reply; if not, it responds with a reply rejecting that recipient (but not the whole mail transaction). The SMTP-sender and SMTP-receiver may negotiate several recipients. When the recipients have been negotiated the SMTP-sender sends the mail data, terminating with a special sequence. If the SMTP-receiver successfully processes the mail data it responds with an OK reply. The dialog is purposely lock-step, one-at-a-time.

The SMTP provides mechanisms for the transmission of mail; directly from the sending user's host to the receiving user's host when the two host are connected to the same transport service, or via one or more relay SMTP-servers when the source and destination hosts are not connected to the same transport service.

To be able to provide the relay capability the SMTP-server must be supplied with the name of the ultimate destination host as well as the destination mailbox name.

The argument to the MAIL command is a reverse-path, which specifies who the mail is from. The argument to the RCPT command is a forward-path, which specifies who the mail is to. The forward-path is a source route, while the reverse-path is a return route (which may be used to return a message to the sender when an error occurs with a relayed message).

When the same message is sent to multiple recipients the SMTP encourages the transmission of only one copy of the data for all the recipients at the same destination host.

The mail commands and replies have a rigid syntax. Replies also have a numeric code. In the following, examples appear which use actual commands and replies. The complete lists of commands and replies appears in Section 4 on specifications.

Commands and replies are not case sensitive. That is, a command or reply word may be upper case, lower case, or any mixture of upper and lower case. Note that this is not true of mailbox user names. For some hosts the user name is case sensitive, and SMTP implementations must take case to preserve the case of user names as they appear in mailbox arguments. Host names are not case sensitive.

Commands and replies are composed of characters from the ASCII character set [1]. When the transport service provides an 8-bit byte (octet) transmission channel, each 7-bit character is transmitted right justified in an octet with the high order bit cleared to zero.

When specifying the general form of a command or reply, an argument (or special symbol) will be denoted by a meta-linguistic variable (or constant), for example, "<string>" or "<reverse-path>". Here the angle brackets indicate these are meta-linguistic variables. However, some arguments use the angle brackets literally. For example, an actual reverse-path is enclosed in angle brackets, i.e., "<John.Smith@USC-ISI.ARPA>" is an instance of (the angle brackets are actually transmitted in the command or reply).

RFC 821 MAIL Procedure

There are three steps to SMTP mail transactions. The transaction is started with a MAIL command which gives the sender identification. A series of one or more RCPT commands follows giving the receiver information. Then a DATA command gives the mail data. And finally, the end of mail data indicator confirms the transaction.

The first step in the procedure is the MAIL command. The <reverse-path> contains the source mailbox.

MAIL <SP> FROM: <reverse-path> CRLF>

This command tells the SMTP-receiver that a new mail transaction is starting and to reset all its state tables and buffers, including any recipients or mail data. It gives the reverse-path which can be used to report errors. If accepted, the receiver-SMTP returns a 250 OK reply.

The <reverse-path> can contain more than just a mailbox. The <reverse-path> is a reverse source routing list of hosts and source mailbox. The first host in the <reverse-path> should be the host sending this command.

The second step in the procedure is the RCPT command.

RCPT <SP> TO: <forward-path> <CRLF>

This command gives a forward-path identifying one recipient. If accepted, the receiver-SMTP returns a 250 OK reply, and
stores the forward-path. If the recipient is unknown the receiver-SMTP returns a 550 Failure reply. This second step of the procedure can be repeated any number of times.

The <forward-path> can contain more than just a mailbox. The <forward-path> is a source routing list of hosts and the destination mailbox. The first host in the <forward-path> should be the host receiving this command.

The third step in the procedure is the DATA command.

DATA <CRLF>

If accepted, the receiver-SMTP returns a 354 Intermediate reply and considers all succeeding lines to be the message text.
When the end of text is received and stored the SMTP-receiver sends a 250 OK reply.

Since the mail data is sent on the transmission channel the end of the mail data must be indicated so that the command and
reply dialog can be resumed. SMTP indicates the end of the mail data by sending a line containing only a period. A transparency procedure is used to prevent this from interfering with the user's text.

Please note that the mail data includes the memo header items such as Date, Subject, To, Cc, From [2].

The end of mail data indicator also confirms the mail transaction and tells the receiver-SMTP to now process the stored recipients and mail data. If accepted, the receiver-SMTP returns a 250 OK reply. The DATA command should fail only if the mail transaction was incomplete (for example, no recipients), or if resources are not available.

The above procedure is an example of a mail transaction. These commands must be used only in the order discussed above. The example below illustrates the use of these commands in a mail transaction.

Example of the SMTP Procedure
This SMTP example shows mail sent by Smith at host Alpha.ARPA, to Jones, Green, and Brown at host Beta.ARPA. Here we assume that host Alpha contacts host Beta directly.

S: MAIL FROM: <Smith@Alpha.ARPA>
R: 250 OK

S: RCPT TO: <Jones@Beta.ARPA>
R: 250 OK

S: RCPT TO: <Green@Beta.ARPA>
R: 550 No such user here

S: RCPT TO: <Brown@Beta.ARPA>
R: 250 OK

S: DATA
R: 354 Start mail input; end with .
S: Blah blah blah...
S: ...etc. etc. etc.
S: .
R: 250 OK

The mail has now been accepted for Jones and Brown. Green did not have a mailbox at host Beta.

RFC 821 Forwarding Procedure

There are some cases where the destination information in the <forward-path> is incorrect, but the receiver-SMTP knows the correct destination. In such cases, one of the following replies should be used to allow the sender to contact the correct destination.

251 User not local; will forward to <forward-path>

This reply indicates that the receiver-SMTP knows the user's mailbox is on another host and indicates the correct forward-path to use in the future. Note that either the host or user or both may be different. The receiver takes responsibility for delivering the message.

551 User not local; please try <forward-path>

This reply indicates that the receiver-SMTP knows the user's mailbox is on another host and indicates the correct forward-path to use. Note that either the host or user or both may be different. The receiver refuses to accept mail for this user, and the sender must either redirect the mail according to the information provided or return an error response to the originating user.

Example of Forwarding
Either

S: RCPT TO: <Postel@USC-ISI.ARPA>
R: 251 User not local; will forward to <Postel@USC-ISIF.ARPA>

Or

S: RCPT TO: <Paul@USC-ISIB.ARPA>
R: 551 User not local; please try <Mockapetris@USC-ISIF.ARPA>

RFC 821 Main SMTP Commands

The SMTP commands define the mail transfer or the mail system function requested by the user. SMTP commands are character strings terminated by . The command codes themselves are alphabetic characters terminated by if parameters follow and otherwise. The syntax of mailboxes must conform to receiver site conventions. The SMTP commands are discussed below. The SMTP replies are discussed in the Section 4.2.

A mail transaction involves several data objects which are communicated as arguments to different commands. The reverse-path is the argument of the MAIL command, the forward-path is the argument of the RCPT command, and the mail data is the argument of the DATA command. These arguments or data objects must be transmitted and held pending the confirmation communicated by the end of mail data indication which finalizes the transaction. The model for this is that distinct buffers are provided to hold the types of data objects, that is, there is a reverse-path buffer, a forward-path buffer, and a mail data buffer. Specific commands cause information to be appended to a specific buffer, or cause one or more buffers to be cleared.

HELLO (HELO)

MAIL (MAIL)

RECIPIENT (RCPT)

DATA (DATA)

RESET (RSET)

QUIT (QUIT)

There are restrictions on the order in which these command may be used.

The first command in a session must be the HELO command. The HELO command may be used later in a session as well. If the HELO command argument is not acceptable a 501 failure reply must be returned and the receiver-SMTP must stay in the same state.

The MAIL command begins a mail transaction. Once started a mail transaction consists of one of the transaction beginning commands, one or more RCPT commands, and a DATA command, in that order. A mail transaction may be aborted by the RSET command. There may be zero or more transactions in a session.

If the transaction beginning command argument is not acceptable a 501 failure reply must be returned and the receiver-SMTP must stay in the same state. If the commands in a transaction are out of order a 503 failure reply must be returned and the receiver-SMTP must stay in the same state.

The last command in a session must be the QUIT command. The QUIT command can not be used at any other time in a session.

This command is used to identify the sender-SMTP to the receiver-SMTP. The argument field contains the host name of the sender-SMTP.

The receiver-SMTP identifies itself to the sender-SMTP in the connection greeting reply, and in the response to this command.

This command and an OK reply to it confirm that both the sender-SMTP and the receiver-SMTP are in the initial state, that is, there is no transaction in progress and all state tables and buffers are cleared.

This command is used to initiate a mail transaction in which the mail data is delivered to one or more mailboxes. The argument field contains a reverse-path.

The reverse-path consists of an optional list of hosts and the sender mailbox. When the list of hosts is present, it is a "reverse" source route and indicates that the mail was relayed through each host on the list (the first host in the list was the most recent relay). This list is used as a source route to return non-delivery notices to the sender. As each relay host adds itself to the beginning of the list, it must use its name as known in the IPCE to which it is relaying the mail rather than the IPCE from which the mail came (if they are different). In some types of error reporting messages (for example, undeliverable mail notifications) the reverse-path may be null.

This command clears the reverse-path buffer, the forward-path buffer, and the mail data buffer; and inserts the reverse-path information from this command into the reverse-path buffer.

This command is used to identify an individual recipient of the mail data; multiple recipients are specified by multiple use of this command.

The forward-path consists of an optional list of hosts and a required destination mailbox. When the list of hosts is present, it is a source route and indicates that the mail must be relayed to the next host on the list. If the receiver-SMTP does not implement the relay function it may user the same reply it would for an unknown local user (550).

When mail is relayed, the relay host must remove itself from the beginning forward-path and put itself at the beginning of the reverse-path. When mail reaches its ultimate destination (the forward-path contains only a destination mailbox), the receiver-SMTP inserts it into the destination mailbox in accordance with its host mail conventions.

For example, mail received at relay host A with arguments

```
FROM: <USERX@HOSTY.ARPA>
TO: <@HOSTB.ARPA:USERC@HOSTD.ARPA>
```

will be relayed on to host B with arguments

```
FROM: <@HOSTA.ARPA:USERX@HOSTY.ARPA>
TO: <@HOSTB.ARPA:USERC@HOSTD.ARPA>
```

This command causes its forward-path argument to be appended to the forward-path buffer.

DATA (DATA)

The receiver treats the lines following the command as mail data from the sender. This command causes the mail data from this command to be appended to the mail data buffer. The mail data may contain any of the 128 ASCII character codes.

The mail data is terminated by a line containing only a period, that is the character sequence "<CRLF>.<CRLF>". This is the end of mail data indication.

The end of mail data indication requires that the receiver must now process the stored mail transaction information. This processing consumes the information in the reverse-path buffer, the forward-path buffer, and the mail data buffer, and on the completion of this command these buffers are cleared. If the processing is successful the receiver must send an OK reply. If the processing fails completely the receiver must send a failure reply.

When the receiver-SMTP accepts a message either for relaying or for final delivery it inserts at the beginning of the mail data a time stamp line. The time stamp line indicates the identity of the host that sent the message, and the identity of the host that received the message (and is inserting this time stamp), and the date and time the message was received. Relayed messages will have multiple time stamp lines.

When the receiver-SMTP makes the "final delivery" of a message it inserts at the beginning of the mail data a return path line. The return path line preserves the information in the from the MAIL command. Here, final delivery means the message leaves the SMTP world. Normally, this would mean it has been delivered to the destination user, but in some cases it may be further processed and transmitted by another mail system.

It is possible for the mailbox in the return path be different from the actual sender's mailbox, for example, if error responses are to be delivered a special error handling mailbox rather than the message senders.

The preceding two paragraphs imply that the final mail data will begin with a return path line, followed by one or more time stamp lines. These lines will be followed by the mail data header and body.

Special mention is needed of the response and further action required when the processing following the end of mail data indication is partially successful. This could arise if after accepting several recipients and the mail data, the receiver-SMTP finds that the mail data can be successfully delivered to some of the recipients, but it cannot be to others (for example, due to mailbox space allocation problems). In such a situation, the response to the DATA command must be an OK reply. But, the receiver-SMTP must compose and send an "undeliverable mail" notification message to the originator of the message. Either a single notification which lists all of the recipients that failed to get the message, or separate notification messages must be sent for each failed recipient (see Example 7). All undeliverable mail notification messages are sent using the MAIL command.

This command specifies that the current mail transaction is to be aborted. Any stored sender, recipients, and mail data must be discarded, and all buffers and state tables cleared. The receiver must send an OK reply.

This command specifies that the receiver must send an OK reply, and then close the transmission channel.

The receiver should not close the transmission channel until it receives and replies to a QUIT command (even if there was an error). The sender should not close the transmission channel until it send a QUIT command and receives the reply (even if there was an error response to a previous command). If the connection is closed prematurely the receiver should act as if a RSET command had been received (canceling any pending transaction, but not undoing any previously completed transaction), the sender should act as if the command or transaction in progress had received a temporary error (4xx).

RFC 821 SMTP Replies

Replies to SMTP commands are devised to ensure the synchronization of requests and actions in the process of mail transfer, and to guarantee that the sender-SMTP always knows the state of the receiver-SMTP. Every command must generate exactly one reply.

The details of the command-reply sequence are made explicit in Section 5.3 on Sequencing and Section 5.4 State Diagrams.

An SMTP reply consists of a three digit number (transmitted as three alphanumeric characters) followed by some text. The number is intended for use by automata to determine what state to enter next; the text is meant for the human user. It is intended that the three digits contain enough encoded information that the sender-SMTP need not examine the text and may either discard it or pass it on to the user, as appropriate. In particular, the text may be receiver-dependent and context dependent, so there are likely to be varying texts for each reply code. A discussion of the theory of reply codes is given in Appendix E. Formally, a reply is defined to be the sequence: a three-digit code, , one line of text, and , or a multi-line reply (as defined in Appendix E). Only the EXPN and HELP commands are expected to result in multi-line replies in normal circumstances, however multi-line replies are allowed for any command.

SMTP Reply Codes

Max Messages

This option allows you to specify the maximum number of incoming messages on a per user or per IP Address basis in any 24 hour period. Any messages exceeding this number should be temporarily rejected with a "450" message.

Maximum Messages Per User

User
Enter the user name that you would like to set a maximum number of messages per day for. The user name should be entered along with the domain name, i.e user@domain.name.

Messages
Enter the maximum number of message this user is allowed to receive in any 24 hour period.

Reject With
Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details. The rejection message used should use an SMTP Reply Code showing a temporary failure so as to ensure that messages do not fail completely.

Maximum Messages Per IP Address

IP Address
Enter the IP Address that you would like to set a maximum number of messages per day for.

Messages
Enter the maximum number of message to be accepted from a particular IP Address in any 24 hour period.

Reject With
Any mail that is rejected should tell the sending server why it is being rejected otherwise they will continue to try to send it to you. Rejection messages should always be preceded with a 3 digit code indicating the reason for refusal. See SMTP Reply Codes for further details. The rejection message used should use an SMTP Reply Code showing a temporary failure so as to ensure that messages do not fail completely.

Log Entry
A typical log entry might look as follows:

SPAM	24 Sep 97	16:37:41	H	3300	13	Send Response	user@another.co.uk [196.198.250.250] -> "450 Too today"
SPAM	24 Sep 97	16:37:41	H	3299	14	Send Response	user@another.co.uk user@mail.net -> "450 Too mar

Redirect File Example

Consider the redirect list:

Action	From	To	Result	Rule
Redirect	*@spammer.com	*	NULL	1
Redirect	postmaster@*	*	postmaster	2
Respond	194.194.194.194/12	*	550 Domain banned	3

Rule 1
If a message arrived from joe@spammer.com, it would be redirected to NULL and lost.

Rule 2
A message from postmaster@spammer.com would be redirected to the local postmaster. Note that although Rule 1 applies as well as Rule 2, the last rule is the one that is executed.

Rule 3
A message from a mail server at 194.194.199.241 will cause the response "550 Domain banned" to be returned after the RCPT clause.

If Rules 2 & 3 were swapped around, a message from the postmaster at domain 194.194.199.241 would be accepted. We recommend that you always allow mail from a postmaster so that if a domain has been incorrectly banned, the problem can be resolved.

Restricted Word File Example
If the phrase was "me", messages with words like "message", "meat", etc., would be foiled. Adding a space either side helps to reduce the liklihood email is incorrectly filtered.