

[MS-XJRNL]: Journal Record Message Format Protocol Specification

Intellectual Property Rights Notice for Protocol Documentation

- **Copyrights.** This protocol documentation is covered by Microsoft copyrights. Regardless of any other terms that are contained in the terms of use for the Microsoft website that hosts this documentation, you may make copies of it in order to develop implementations of the protocols, and may distribute portions of it in your implementations of the protocols or your documentation as necessary to properly document the implementation. This permission also applies to any documents that are referenced in the protocol documentation.
- **No Trade Secrets.** Microsoft does not claim any trade secret rights in this documentation.
- **Patents.** Microsoft has patents that may cover your implementations of the protocols. Neither this notice nor Microsoft's delivery of the documentation grants any licenses under those or any other Microsoft patents. However, the protocols may be covered by Microsoft's Open Specification Promise (available here: <http://www.microsoft.com/interop/osp>). If you would prefer a written license, or if the protocols are not covered by the OSP, patent licenses are available by contacting protocol@microsoft.com.
- **Trademarks.** The names of companies and products contained in this documentation may be covered by trademarks or similar intellectual property rights. This notice does not grant any licenses under those rights.

Reservation of Rights. All other rights are reserved, and this notice does not grant any rights other than specifically described above, whether by implication, estoppel, or otherwise.

Tools. This protocol documentation is intended for use in conjunction with publicly available standard specifications and network programming art, and assumes that the reader either is familiar with the aforementioned material or has immediate access to it. A protocol specification does not require the use of Microsoft programming tools or programming environments in order for you to develop an implementation. If you have access to Microsoft programming tools and environments you are free to take advantage of them.

Revision Summary			
Author	Date	Version	Comments
Microsoft Corporation	April 4, 2008	0.1	Initial Availability.
Microsoft Corporation	April 25, 2008	0.2	Revised and updated property names and other technical content.
Microsoft Corporation	June 27, 2008	1.0	Initial Release.
Microsoft Corporation	August 6, 2008	1.01	Revised and edited technical content.
Microsoft Corporation	September 3, 2008	1.02	Revised and edited technical content.

Table of Contents

1	<i>Introduction</i>	3
1.1	Glossary	3
1.2	References	4
1.2.1	Normative References	4
1.2.2	Informative References	5
1.3	Structure Overview	5
1.3.1	Body Text of the Journal-Report Message	5
1.3.2	Original-Message	5
1.4	Relationship to Protocols and Other Structures	5
1.5	Applicability Statement.....	6
1.6	Versioning and Localization.....	6
1.7	Vendor-Extensible Fields	6
2	<i>Structures</i>	7
2.1	Envelope-Part	7
2.1.1	Sender	7
2.1.2	Subject.....	7
2.1.3	Message-ID.....	7
2.1.4	Recipient-specification.....	7
2.1.4.1	recipient-p2-type.....	8
2.1.4.2	forward-path.....	8
2.1.4.3	redirection-type.....	8
2.1.4.3.1	Expanded	8
2.1.4.3.2	Forwarded.....	8
2.1.4.4	original-forward-path	9
2.1.5	label.....	9
2.2	Original-Message-Part	9
3	<i>Structure Examples</i>	9
4	<i>Security Considerations</i>	10
5	<i>Appendix A: Office/Exchange Behavior</i>	10
	<i>Index</i>	12

1 Introduction

Journal record messages are e-mail messages that are generated by the server to capture and report information about messages sent to and from users of the e-mail system. For background information about how journaling works, see [MSFT-WPXTJ].

This document specifies an extension to [RFC2045] and [RFC2046] to accommodate journal record messages.

1.1 Glossary

The following terms are defined in [MS-OXGLOS]:

Augmented Backus-Naur Form (ABNF)
binary large object (BLOB)
body part
character set
distribution list
journal
Journal-Report
message ID (MID)
MIME
MIME Content-Type
MIME message
Original-Message
Simple Mail Transfer Protocol (SMTP)
recipient

The following terms are specific to this document:

body: The term contents of a **body part** or an entire message (containing several **body parts**), as specified in [RFC2045] section 2.6.

Content-Transfer-Encoding: A header field that defines a way of **encoding** non-ASCII character data as ASCII data. This is specified in [RFC2045] section 6.

encoding: The process of specifying a **Content-Transfer-Encoding** to transform character data from one form to another.

Envelope-Part: The portion of a **Journal-Report** containing metadata about the **journalled** message in machine-readable form.

message/rfc822: The **MIME Content-Type** of an e-mail message that is embedded within another e-mail message. The type is expressed via the well known **MIME Content-Type** header field. The **message/rfc822** type is specified in [RFC2046] section 5.2.

MIME attachment: A **body part** within a **MIME message**, for example, an e-mail message or file that is attached to an e-mail message.

Original-Message-Part: The portion of a **Journal-Report** that captures the contents of the **Original-Message** that was **journalled**. The **Original-Message-Part** contains the entire data of the **Original-Message**. The **Original-Message-Part** is an opaque **binary large object (BLOB)**.

recipient forwarding: A feature provided by many e-mail delivery systems, where a message sent to one **recipient** e-mail address is instead redirected to a different e-mail address, called the "forwarded address". E-mail software typically provides methods to configure which specific e-mail addresses are enabled for **recipient forwarding**. Some e-mail systems allow configuring whether the message is entirely redirected to the forwarded address with no copy going to the original e-mail address, or whether the original e-mail address receives a copy of the message *in addition to* the forwarded address.

MAY, SHOULD, MUST, SHOULD NOT, MUST NOT: These terms (in all caps) are used as described in [RFC2119]. All statements of optional behavior use either MAY, SHOULD, or SHOULD NOT.

1.2 References

1.2.1 Normative References

[MS-OXGLOS] Microsoft Corporation, "Exchange Server Protocols Master Glossary", June 2008.

[MS-OXMSG] Microsoft Corporation, ".MSG File Format Specification", June 2008.

[RFC2045] Freed, N., et al., "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996, <http://www.ietf.org/rfc/rfc2045.txt>.

[RFC2046] Freed, N. and Borenstein, N., "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, November 1996, <http://www.ietf.org/rfc/rfc2046.txt>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <http://www.ietf.org/rfc/rfc2119.txt>.

[RFC2821] Klensin, J., "Simple Mail Transfer Protocol", RFC 2821, April 2001, <http://www.ietf.org/rfc/rfc2821.txt>.

[RFC2822] Resnick, P., Ed., "Internet Message Format", RFC 2822, April 2001, <http://www.ietf.org/rfc/rfc2822.txt>.

[RFC4234] Crocker, D., Ed. and Overell, P., "Augmented BNF for Syntax Specifications: ABNF", RFC 4234, October 2005, <http://www.ietf.org/rfc/rfc4234.txt>.

1.2.2 Informative References

[MSFT-WPXTJ] Microsoft Corporation, "White Paper: Exchange 2007 Transport Journaling", September 2007, <http://go.microsoft.com/fwlink/?LinkId=111316>.

1.3 Structure Overview

Journal-Reports are **MIME messages** produced by the server that capture information about other ordinary (non-Journal-Report) messages sent through the server. These other messages are referred to here as **Original-Messages**.

1.3.1 Body Text of the Journal-Report Message

The **body** text of the Journal-Report message lists the e-mail addresses of the sender and **recipients** of the message, the subject, the Internet **message ID (MID)**, and certain other metadata about the Original-Message. The **body** text is referred to as the **Envelope-Part** of the Journal-Report. This document formally specifies the structure of the Envelope-Part.

1.3.2 Original-Message

The Original-Message is attached as a **MIME attachment** to the Envelope-Part. This is referred to as the **Original-Message-Part** of the Journal-Report. How the Original-Message-Part is attached to the Envelope-Part is fully specified in [RFC2045] and [RFC2046].

1.4 Relationship to Protocols and Other Structures

The Envelope-Part that is documented here is a substructure within the larger structure of the Journal-Report MIME message. The Journal-Report structure is a MIME message and conforms to [RFC2045] and [RFC2046].

[RFC2045] specifies how messages with a **MIME Content-Type** of **message/rfc822** might be nested recursively as attachments. The outermost **message/rfc822 body part** of the Journal-Report contains the Envelope-Part as the body.

The Envelope-Part is encoded using the mechanisms prescribed in [RFC2045], such as the Content-Transfer-Encoding mechanism, which specifies details such as the **character set** and encoding used for the data in the Envelope-Part. This document specifies the syntax of the Envelope-Part prior to any **MIME encoding** being applied.

To illustrate the concept of encoding: the Envelope-Part could contain textual data outside the ASCII character set range if the Original-Message had any non-English text. The MIME

message is in many circumstances constrained to contain only data in the ASCII character range, so the MIME standards define methods of encoding non-ASCII data using mechanisms such as "base64" and "quote-printable encoding." Thus, if an encoding mechanism was specified for the Envelope-Part using standardized MIME conventions, decoding would be done before processing the Envelope-Part structure.

The terms **Content-Transfer-Encoding**, character set, and encoding are specified in [RFC2045]. The mechanism for decoding the Envelope-Part is specified in [RFC2045] section 6.

Figure 1 shows how the Envelope-Part substructure is placed in relation to the various other substructures in the Journal-Report MIME message.

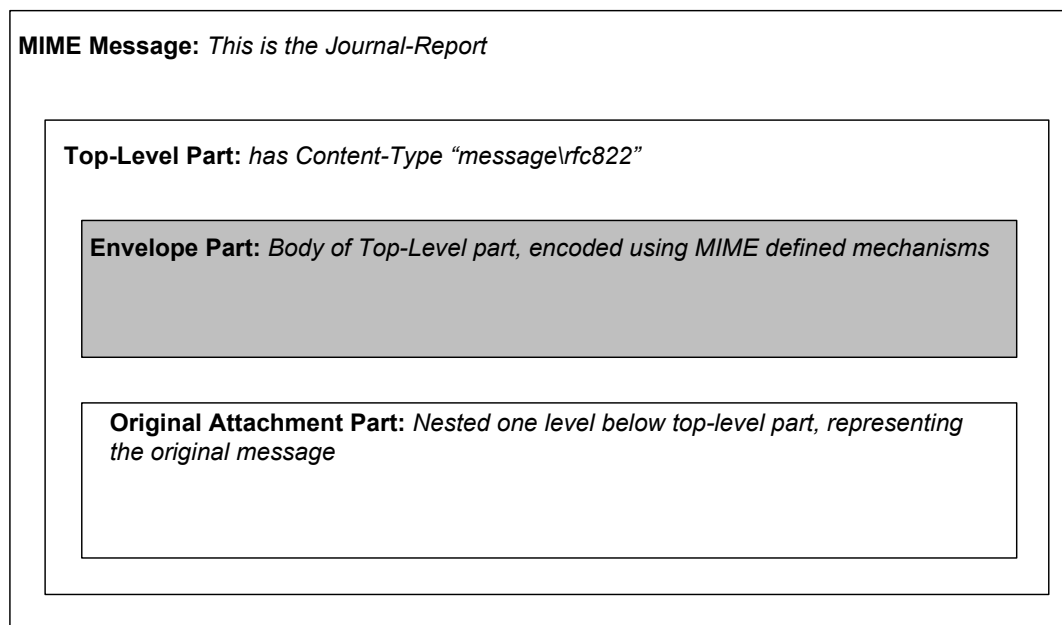


Figure 1: MIME structure of a Journal-Report message

1.5 Applicability Statement

Applications can use this protocol to create and consume journal record messages.

1.6 Versioning and Localization

None.

1.7 Vendor-Extensible Fields

None.

2 Structures

Journal-Reports extend the structures defined in [RFC2045] and [RFC2046] by defining a substructure called the Envelope-Part, which is embedded within the MIME message.

2.1 Envelope-Part

The **Augmented Backus-Naur Form (ABNF)** notation [RFC4234] is used to specify the format of the Envelope-Part. The format is a series of field/value pairs on CRLF terminated lines.

```
<Envelope-Part> =  
  
    "Sender:" SP <reverse-path> CRLF  
  
    "Subject:" SP <subject-string> CRLF  
  
    "Message-ID:" SP <msg-id> CRLF  
  
    1*<recipient-specification>  
  
    0*1<label>
```

2.1.1 Sender

The **reverse-path** field MUST be a **Simple Mail Transfer Protocol (SMTP)** e-mail address as specified in [RFC2821]. This field MUST be set to the sender of the Original-Message.

2.1.2 Subject

The **subject-string** field MUST contain the data from the "Subject" header of the **Original-Message**. This header is specified in [RFC2822].

The **subject-string** can consist of characters outside the ASCII character-set range as specified in [RFC2045] and [RFC2046]. The **MIME Content-Type** header of the respective **body part** in which **Envelope-Part** is embedded MUST specify the **character set** to use to interpret the **subject-string** in accordance with the MIME specifications [RFC2045] and [RFC2046].

2.1.3 Message-ID

The **msg-id** field MUST contain the value of the SMTP "Message-ID" header of the **Original-Message**. This header is specified in [RFC2822] section 3.6.4.

2.1.4 Recipient-specification

One or more occurrences of **recipient-specification** give information about the recipients of the Original-Message that was journaled.

The following provides details about **recipient-specification**:

```
<recipient-specification> =  
    <recipient-p2-type> ":" SP <forward-path>  
    [<redirection-type> ":" SP <original-forward-path>] CRLF  
    <recipient-p2-type> = "Bcc" / "To" / "Cc" / "Recipient"  
    <redirection-type> = "Expanded" / "Forwarded"
```

2.1.4.1 recipient-p2-type

This field MAY be "Bcc," "To," or "Cc," denoting respectively that in the Original-Message the recipient listed in **forward-path** was addressed as a Bcc, To, or Cc recipient when the user sent the message using a mail client. If the server is not able to determine how the user sent the message in the mail client, the **recipient-p2-type** is set to "Recipient."

2.1.4.2 forward-path

This field MUST be an SMTP e-mail address as specified in [RFC2821]. This field MUST be set to a recipient of the Original-Message.

2.1.4.3 redirection-type

This field MUST be set to either "Expanded" or "Forwarded."

2.1.4.3.1 Expanded

A message sent to a **distribution list** is redirected instead to the users and distribution groups contained within the distribution list. This process is known as "expansion" and the distribution group is said to have been "expanded." The expansion is repeated for the nested distribution groups until all distribution groups have been expanded to ordinary users. The **redirection-type** element, when set to "Expanded," denotes that the sender of the message sent it originally to **original-forward-path**, which was a distribution group. The server then expanded **original-forward-path** to one or more recipients (perhaps expanding nested recipients repeatedly until all recipients were non-distribution group, ordinary recipients). Each of these expanded recipients is listed in the **forward-path** field of a **recipient-specification**.

2.1.4.3.2 Forwarded

A **redirection-type** of "Forwarded" indicates that the recipient indicated by **original-forward-path** was configured for **recipient forwarding**. The message was forwarded to the recipient indicated by **forward-path**.

2.1.4.4 original-forward-path

This field MUST be an SMTP e-mail address as specified in [RFC2821]. The **original-forward-path** field is the SMTP address of the recipient that was redirected to **forward-path** by the server. Although syntactically identical to **forward-path**, it has different semantics. The appearance of **redirection-type** and **original-forward-path** in a **recipient-specification** indicates that originally the message that was journaled was sent to the recipient with address **original-forward-path** and the server did one of the following:

1. Changed the recipient address to **forward-path**.
2. Added a new recipient with the address **forward-path** due to the **original-forward-path** recipient being present.

If the server does not have information about what type of forwarding or expansion was done on the recipient, the **redirection-type** and **original-forward-path** elements are omitted.

2.1.5 label

This optional field SHOULD contain a text value, the contents of which are implementation-specific.

The following provides detail about **label**:

```
<label> = "Label:" SP 1*255CHAR CRLF
```

2.2 Original-Message-Part

This data SHOULD be one of the following:

- A MIME attachment of type message/rfc822.
- A .MSG file as specified in [MS-OXMSG].

3 Structure Examples

The following is an example of the Envelope-Part of Journal-Report, followed by an explanation of the various syntax elements (note that the line numbers are not present in the actual Envelope-Part, but are shown here so the structure can be discussed line by line):

```
1.      Sender: sender@contoso.com
2.      Subject: Sample Message
3.      Message-ID: <12345@contoso.com>
4.      To: dl-to-member1@contoso.com, Expanded: dl-to@contoso.com
5.      To: dl-to-member2@contoso.com, Expanded: dl-to@contoso.com
6.      Cc: fwd@contoso.com, Forwarded: user@contoso.com
7.      Bcc: dl-bcc-member@contoso.com, Expansion: dl-
        bcc@contoso.com
8.      Bcc: fwd@contoso.com, Forwarded: user@contoso.com
```

9. Recipient: user-unk@contoso.com
1. The sender of the Original-Message was "sender@contoso.com."
 2. The subject of the Original-Message was "Sample Message."
 3. The message-ID of the Original-Message was "12345@contoso.com."
 4. The Original-Message was sent by the mail client as To: to dl-to@contoso.com, which is a distribution list that was expanded to dl-to-member1@contoso.com and dl-to-member2@contoso.com (captured in Line 5) by the server.
 5. See (4).
 6. The Original-Message was sent to user@contoso.com as Cc:, which was changed by the e-mail server to fwd@contoso.com because recipient forwarding was configured. User@contoso.com did not receive a copy of the message because there is no **recipient-specification** where user@contoso.com was listed in the **forward-path**.
 7. The Original-Message was sent by the mail client as Bcc: to dl-bcc@contoso.com, which is a distribution list that was expanded to dl-bcc-member@contoso.com and possibly other recipients.
 8. The Original-Message was sent to user@contoso.com as Bcc:, which was rewritten by the e-mail server to fwd@contoso.com because recipient forwarding was configured. User@contoso.com did not receive a copy of the message because there is no **recipient-specification** where user@contoso.com was listed in the **forward-path**.
 9. Finally, there is no information about whether user-unk@contoso.com was sent the Original-Message as a To, Cc, or Bcc recipient. It is also not known whether this recipient got the message due to distribution list expansion, recipient forwarding, or was directly addressed by the sender. The "Recipient:" element indicates that the server was only able to capture that user-unk@contoso.com was a recipient of the message and no further recipient metadata was available.

4 Security Considerations

This format does not implement or concern itself with security, but relies on the underlying e-mail transport software and e-mail storage software to provide security as applicable.

5 Appendix A: Office/Exchange Behavior

The information in this specification is applicable to the following versions of Office/Exchange:

- Office 2003 with Service Pack 3 applied
- Exchange 2003 with Service Pack 2 applied

- Office 2007 with Service Pack 1 applied
- Exchange 2007 with Service Pack 1 applied

Exceptions, if any, are noted below. Unless otherwise specified, any statement of optional behavior in this specification prescribed using the terms SHOULD or SHOULD NOT implies Office/Exchange behavior in accordance with the SHOULD or SHOULD NOT prescription. Unless otherwise specified, the term MAY implies Office/Exchange does not follow the prescription.

Index

- Applicability statement, 6
- Envelope-Part structure, 7
- Fields, vendor-extensible, 6
- Glossary, 3
- Informative references, 5
- Introduction, 3
- Normative references, 4
- Office/Exchange behavior, 10
- Original-Message-Part structure, 9
- Overview, 5
- References, 4
 - Informative references, 5
 - Normative references, 4
- Relationship to protocols and other structures, 5
- Security considerations, 10
- Structure examples, 9
- Structures, 7
 - Envelope-Part, 7
 - Original-Message-Part, 9
- Vendor-extensible fields, 6
- Versioning and localization, 6