

Notes from Support: Conversion services of the SMTP MTA

by Bret Swedeen

[Editor's note: This article resides in "Iris Today", the technical Webzine located on the <http://www.notes.net> Web site produced by Iris Associates, the developers of Domino and Notes.]

The Under the Microscope column is brought to you by Lotus Support, and features in-depth technical articles that explore how a particular Domino/Notes feature works. This month's article is the second article in a three-part series on the SMTP MTA, and explores the conversion services. The [first article](#) focused on the transport services of the SMTP MTA. The third, final article will explore the Delivery Report Task (DRT) and Add-in task.

The [first article](#) in this three-part series introduced the transport services of the SMTP Message Transfer Agent (MTA). These services handle the sending and receiving of e-mail between Notes and the Internet. If you have a solid understanding of how the MTA's transport services work, you already understand a large portion of how the MTA works. Of course, as you can already guess, message delivery isn't the only thing the MTA does.

The other important job of the SMTP MTA is the conversion of messages from Notes format to Internet format, and vice versa. Since Notes and the Internet structure a message differently, the MTA must convert a message from one format to another before delivering it to its destination.

This article explores the individual components that make message conversion possible, and explains how these components work together (however, if you want an introductory look at the SMTP MTA, see the *Iris Today* article "[Internet messaging with the SMTP MTA](#)"). Also, in keeping with the first article, a few hands-on exercises are included to further extend your understanding of the conversion services, and to help you put them *Under the Microscope*.

Note: The "Internet format" is used throughout this article as an abbreviation for an Internet-standard message format based on RFC documents.

Message anatomy 101

Understanding the MTA's message conversion services starts with a basic understanding of Internet and Notes message anatomy.

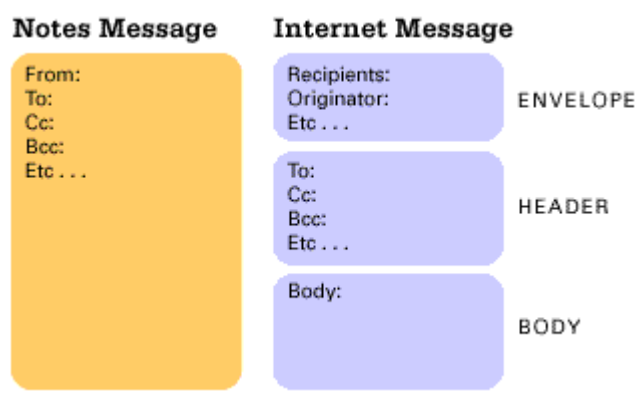
Internet mail messages

The format of an Internet mail message is defined by a few Request for Comments (RFC) documents. These documents describe the various field labels and message structure of a standard Internet electronic mail message. These documents make for some pretty dull and detailed reading. Therefore, for the sake of brevity, here is a quick summary of three primary sections that make up an Internet electronic mail message:

- The envelope: This section contains the information used during an SMTP client and server message exchange session (this part of the message is defined in [RFC 821](#), which also defines the SMTP transport).
- The header: This section contains message information, such as the sender, recipient, delivery date, and subject line (this part of the message is defined in [RFC 822](#)).
- The body: This section contains the actual correspondence part of the message and file attachments. When the MTA is configured for "Simple Text," RFC 822 defines the message body. Any other configuration, however, indicates the message body conforms to Multipurpose Internet Message Extensions (MIME) format, which is defined in [RFCs 2045](#) and [2049](#).

Notes mail messages

The format of a Notes mail message is very similar to an Internet mail message; however, it's not as neatly compartmentalized. A Notes message is basically a document with a list of fields containing important delivery information, such as the sender, recipient, and subject line. The Body field in the message contains the message correspondence and file attachments. Some of the fields included in a Notes message are visible to the user and others are hidden.



Presto chango -- The magical components of conversion

The magic of conversion isn't the work of one special server task or Notes database. Message conversion is a process that involves several individual components, each doing their part along the path of conversion. While some of these components might do more "work" than the others, each piece is a critical part of the whole. Remove one of these components, and message conversion is incomplete.

What follows is a list of these critical components in no particular order:

SMTP.BOX

This component is a Notes database located on the Domino server running the SMTP MTA. This database holds Internet-bound messages waiting for conversion. The Outbound Conversion Handler reads the message from this database, converts it, and places the "Internet-ready" message in the Outbound Work Queue for transmission.

MAIL.BOX

The MAIL.BOX is a Notes database similar in design to the SMTP.BOX. This database holds Notes-bound messages already converted from Internet format to Notes format. The Router (a Domino server task) takes the message from this database for delivery within the Notes environment.

Inbound Work Queue database

This database (SMTPIBWQ.NSF) is used for temporary message storage. An Inbound Session Handler places messages received from the Internet in this database.

Outbound Work Queue database

This database (SMTPOBWQ.NSF) is another temporary storage database. The Outbound Session Controller takes converted messages from this database and hands them over to an Outbound Session Handler for transmission.

SMTP Tables database

This database is used by the Inbound and Outbound Conversion Handlers as a lookup table for character set and file type identification routines in the MTA. This database (SMTPTBLS.NSF) is also a cross-reference table between MIME type/subtypes and file extensions.

MTA Forms database

This database (MTAFORMS.NSF) is used by the Outbound Conversion Handler to correctly render a message. If your Notes environment uses custom letterheads, you must add them to this database to ensure proper message conversion.

Note: You can add each of these special databases to your Notes workspace by choosing File - Database - Open and selecting them from the Open Database dialog box. The only exceptions are the two .BOX databases. To add these databases to your workspace, you need to enter MAIL.BOX or SMTP.BOX in the Filename field of the Open Database dialog box. Keep in mind, that even though you can add these databases to your workspace, Lotus strongly discourages any type of modification beyond what is described in the [Domino Administrator's Guide](#). Additional views or form customizations might disrupt the functional purpose of these databases.



Inbound conversion handler

This task is started by the SMTP add-in task when the MTA loads. The Inbound Conversion Handler appears on the server console as **SMTPMTA imsgcnv**. This handler remains idle until an Inbound Session Handler tells it that there is a message ready for conversion in the Inbound Work Queue. Once notified, this handler moves into action and starts the message conversion process. If there is a problem with the new message notification from an Inbound Session Handler, the Inbound Conversion Handler periodically checks the Inbound Work Queue for new messages.

Outbound conversion handler

This task is also started by the SMTP add-in task when the MTA loads and appears on the server console as **SMTPMTA omsgcnv**. This handler periodically checks the SMTP.BOX for outbound messages. Once a message is found, the handler converts the message to Internet format, places the converted message in the Outbound Work Queue, and informs the Outbound Session Controller that a message is ready for transmission.

The black box of conversion

At first glance, message conversion looks simple. It's true that message conversion in its most basic form involves extracting text from one format and placing it into another. Message conversion in the real world, however, is much more than cut-and-paste. A closer look inside the "black box" of conversion helps reveal the true complexities of message conversion.

Outbound message conversion

The outbound message conversion process begins once the Domino Router places an Internet-bound message in the SMTP.BOX database. The Outbound Message Handler continually checks the SMTP.BOX for new messages. Once a new message is found, the handler creates a document stub for this message in the Outbound Work Queue. This stub holds each piece of the Internet message as the conversion handler creates it. From this point, the conversion handler begins the "real" work of transforming the Notes message into the three parts required by the Internet message format -- the envelope, message header, and message body.

Note: In a busy environment, the Outbound Message Handler can encounter numerous messages when it checks the SMTP.BOX. The handler, however, can only read 100 messages at a time. Rather than read the first 100 messages that arrive in the SMTP.BOX, the handler reads the first 100 messages based on their Notes IDs.

Envelope construction

Message transformation begins with the envelope construction. Since a Notes message doesn't have an envelope, the handler builds one from information contained in the Notes message. The handler derives the envelope's recipients list from the Notes recipients field, which is a combination of the To, cc, and bcc fields. It then adds the addresses one-by-one to the corresponding envelope field.

Note: While the Notes recipients field contains all recipients, the Router removes each Notes recipient from this field as it successfully delivers the message to that recipient. By the time the message reaches the SMTP.BOX, only the Internet addresses remain in the recipients field.

After the handler populates the recipients list, it determines the Next Hop Entity (either the message's destination domain or a mail relay host). This information is used by the MTA transport services when delivering the message. The handler looks at each recipient and determines the Next Hop Entity for that recipient based on the settings in the Foreign Domain document and SMTP Connection document in the Public Address Book.


The final part of envelope construction involves populating the originator field (the person who sent the message). In Notes, however, there is no originator, there is only a Sender (the From field). Therefore, at this point, the conversion handler converts the From field to a legal Internet address, which becomes the message originator. This part of the process is sometimes problematic because legal characters in Notes aren't always legal on the Internet (and vice versa). The handler makes every attempt possible to convert illegal characters; however, once all possibilities are exhausted, message conversion fails and the sender is notified of the failure.

Header construction

After constructing the envelope, the conversion handler creates the message header and populates the RFC 822 header fields with their Notes counterparts. Most of these fields, such as Date and Subject, are direct mappings. Address fields (To, cc, bcc) also directly map to an RFC 822 counterpart; however, each address is checked for Internet legality to ensure proper "reply-to-all" delivery. If an address is not already in Internet format, the handler must first convert the address from Notes format to Internet format. If the address contains illegal characters (by Internet standards), the handler makes every possible attempt to handle these characters. If address conversion fails at this point, however, the message is still delivered. The only problem is that when the recipient sends a "reply-to-all," everyone receives the reply except any addresses that didn't properly convert.

Note: Prior to Release 4.6.1, if there was a problem with address conversion, message conversion failed and the problematic message was returned to the sender.

The final part of header conversion is the management of bcc (blind carbon copy) recipients. The recipients in this field receive a copy of the message, but other recipients don't know it. Therefore, for each bcc recipient, the handler creates a separate message in the Outbound Work Queue. For example, if you have two bcc recipients, one primary and one cc recipient, the conversion handler creates three outbound messages -- one for each bcc recipient, and one for the primary and cc recipients.



Body construction

The final, and most involved, part of message conversion is the creation of the message body. At this final stage, the conversion handler makes an API call that breaks the message into individual Notes items. For example, body text is an item, an in-line graphic is an item, and an attachment is an item. Once the API creates these items, it passes them to Notes Conversion Services (CVS), which creates "SMTP-ready" message items.

Once CVS is finished building the "SMTP-ready" message items, the conversion handler takes these items and builds the message body according to the Message Content setting (Simple Text, Only Lotus Notes Users, Users without Lotus Notes, and Users with and without Lotus Notes) in the Internet Message Transfer Agent (SMTP MTA) section of the Server document. For the moment, let's assume this field is set to "Users without Lotus Notes" so we can see how MIME body content is created.

The conversion handler takes the message items created by CVS and constructs a MIME message. If a message item is identified as an attachment, the conversion handler makes an API call to determine the file type. If the API successfully determines the file type, the conversion handler checks the MTA Tables database and locates the corresponding MIME type/subtype for this file type. If the API cannot determine the file type, the conversion handler uses the file's extension to locate the correct MIME type/subtype in the MTA Tables database. The attachment is then encoded using the setting specified in the Server document and then placed in the message body.

Note: Users can override the default encoding method when they send a message by choosing Actions - Special Options, and then selecting an alternative encoding method -- MIME (Base 64), Quoted Printable, or UUencode.



Now that you've seen how the handler creates MIME body content, let's look at how it constructs the other types of Message Content:

- **Simple Text Only:** For simple text messages, the conversion handler takes the "SMTP-ready" items from CVS and puts them straight into the message body. The handler doesn't need to cross-reference the MTA Tables database because the messages are simple text and not MIME. Attachments in this situation are added to the message as simple text after they've been encoded using UUencode.
- **Only Lotus Notes Users:** For messages intended only for Notes users, the conversion handler takes the "SMTP-ready" items and creates an ENCAP2.OND file. This type of file is a "minimalist" Notes database (no views, other than the default view, no forms, and only one document), which is as small as possible for efficient transport. The handler encodes the file and adds it to the message body as MIME body content with a boundary identifier of "MIME Application / X.Lotus-Encap2."
- **Users with and without Lotus Notes:** For messages intended for both Notes and non-Notes users, the handler creates a MIME Multipart Alternative message body. Essentially, the message contains two message parts. The first part is a standard body content message as described previously, which is meant for users without Lotus Notes. The other message option is an ENCAP2.OND file, which is meant for Lotus Notes users.

Regardless of how the handler constructs the message body, once it converts that message and temporarily stores it in the Outbound Work Queue, the conversion process is complete. From this point forward, the transport services of the SMTP MTA assume responsibility for the message.

Inbound message conversion

The inbound conversion process begins when the Inbound Work Queue database receives a new message. The conversion handler is notified by an Inbound Session Handler through inter-process communication that a new message is ready for conversion. Once notified, the conversion handler creates a message stub for the new message in the MAIL.BOX and then begins converting the three Internet message parts (envelope, header, and body) into a single Notes message.

Envelope conversion

The handler begins converting the envelope by first analyzing each address in the recipients field. If none of the addressees are members of the MTA's Notes domain, the handler does not convert the message. Instead, it places the message in the Outbound Work Queue, where the Outbound Session Controller takes care of delivering the message to its Next Hop Entity (this situation occurs when the MTA is relaying a message to another host as described in [RFC 821](#)). On the other hand, for addresses that do correspond to the MTA's Notes domain, the handler converts them to Notes addresses and then adds the addresses to the Notes recipient fields. Lastly, the handler stores the envelope's originator in the SMTPOriginator field in the Notes message because Notes doesn't have an equivalent to an originator.

Note: Usually the Internet name is stored in the Short Name field in the Person document, which makes the name conversion simple. Sometimes, however, name conversion is not so simple. Therefore, for a detailed explanation of Internet name conversion, refer to Tech Note #145300 "How Does the SMTP MTA Process Addresses?" in the [Lotus Notes Knowledge Base](#).

Header conversion

The handler converts the message header by mapping RFC fields to corresponding Notes fields. This information is also stored in the \$AdditionalHeaders field (if configured as such), which you can see when you choose Actions - Delivery Information. The handler doesn't convert the From field, but does mask characters that Notes can't handle. And finally, the handler analyzes the addresses contained in the TO, CC, and BCC fields. If the address is for a Notes user, it converts the address to its corresponding Notes address and then places it into the appropriate field. Non-Notes addresses are not converted -- it simply places them into the corresponding Notes field.

Note: Some headers contain a large number of addresses that when combined together might approach, or exceed, the Notes 15k summary information limit. When the MTA anticipates such a problem, it starts to truncate the address fields once it hits 13k in order to allow room for the other summary items Notes adds for message delivery. The only downside to this truncating is that if the recipient tries to send a "Reply to All," the delivery will fail for the addresses that were truncated.

Body conversion

The Inbound Conversion Handler starts the body conversion process by analyzing the MIME boundary identifiers. As the handler moves through the message from the top to the bottom, it builds a message structure of the SMTP items contained in the body. If an item is identified as a file, the conversion handler makes an API call to determine the file type. If the API successfully determines the file type, the handler uses that number to cross-reference the MTA Tables database for a file description and the proper extension for the file (if necessary). If the API cannot determine the file type, the handler relies on the MIME type/subtype of the item to properly cross-reference it. In situations where an item has no name, the handler uses the file extension to cross-reference the MTA Tables database.

Once the handler builds the message structure and properly identifies each item, it passes the structure to CVS. At this point, CVS converts the message items into their Notes equivalent and passes them back to the Inbound Conversion Handler which uses them to populate a message stub in the MAIL.BOX. The Router then delivers the message to its recipients, and the recipients view the message using the default database form.



Hands-on work with the SMTP MTA

Technical explanations and pretty-looking diagrams can help illustrate how message conversion fundamentally works. However, the best way to completely understand the conversion services of the SMTP MTA is to do a little hands-on work with the feature itself. The exercises below are designed to help reinforce the concepts presented in this article.

Hands-on environment setup

To complete the exercises, you'll need the following:

- Two computers (one for the Domino server and one for the Notes client)
- Domino server 4.x (installed and configured with TCP/IP)
- SMTP MTA 1.x (installed and configured for manual loading, that is, do not let the Domino server automatically start the MTA)
- A Notes client (installed and configured for a registered Notes user, and communicating with your Domino server via TCP/IP)
- At least one valid Internet e-mail address to which you can send test messages

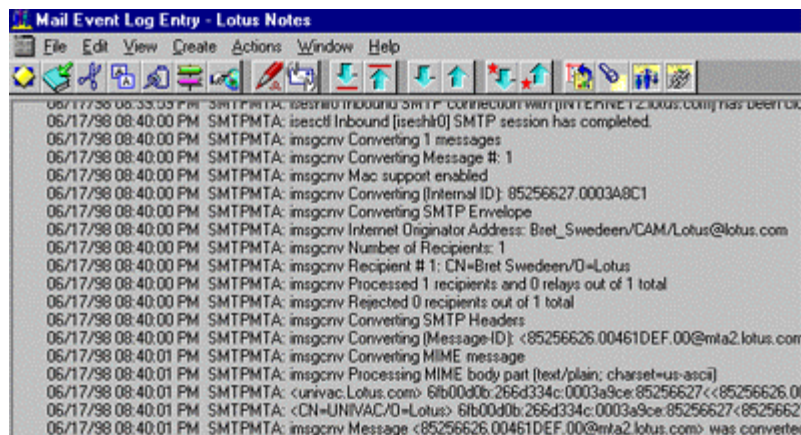
Note: If you are using Domino Release 4.6, the SMTP MTA version number is the same as the Domino release number.

Exercise #1: What the MTA says about message conversion

To see the exact process of message conversion, follow the steps listed below:

1. In the Public Address Book located on this machine, open the Server document for the machine running the SMTP MTA.
2. Edit the Internet Message Transfer Agent (SMTP MTA) section in the Server document and set the Log Level to Verbose (this setting causes the MTA to display detailed informational messages on the server console and in the Notes Log database).
3. Save your changes.
4. Edit the server's NOTES.INI file and add "SMTPMTA_TRACE_LEVEL=0x00000002 (this setting records nearly every action performed by the Outbound Conversion Handler in the Notes Log database).
5. Save your changes.
6. Restart the server.
7. Once the MTA is fully initialized, send one of your Notes users a message from the Internet.
8. Immediately after you send the message, watch the server console.

You should see more informational messages than you ever dreamed were possible (talk about information overload). These messages indicate exactly what the conversion handler did once a new message was received. If the informational messages scrolled by too quickly to read, open the Notes Log database and look under Mail Routing Events. Everything that appeared on the server console is contained in a document under this view.



Exercise #2: Conversion services' scope of responsibility

Based on the description of the outbound conversion process, you might assume that once a message is converted, it's removed from the SMTP.BOX. This assumption is not correct. The message in the SMTP.BOX is the parent document of its converted version (a.k.a. the child document) in the Outbound Work Queue. If the message transport fails, the parent is used to help deliver a message failure report to the sender.

The relationship between the parent document in the SMTP.BOX and the child document in the Outbound Work Queue isn't necessarily about conversion; however, their relationship helps demonstrate the conversion services' scope of responsibility.

1. Choose File - Database - Open and add the SMTP.BOX and Outbound Work Queue databases to your Notes workspace. Remember, you can select the Outbound Work Queue from the list of databases in the dialog box; however, to open the SMTP.BOX, you need to enter SMTP.BOX as the file name.
2. If the SMTP MTA isn't running on your server, enter LOAD SMTPMTA at the server console.
3. Once the MTA is initialized, enter TELL SMTPMTA STOP DRTASK at the server console. This command stops the Delivery Report Task (DRT task), which is responsible for cleaning up the SMTP.BOX and Outbound Work Queue.
4. Send a message to someone on the Internet. Watch the informational messages on the server console until one of the messages indicates the transmission is complete.
5. Switch to the Notes workspace and refresh the view in the SMTP.BOX and the Outbound Work Queue. Notice the parent document of the message you just sent in the SMTP.BOX and the child (a.k.a. converted message) document in the Outbound Work Queue are still listed. Both of these messages remain in their respective databases despite successful conversion and transmission.
6. Switch back to the server console and enter TELL SMTPMTA START DRTASK.
7. Wait a minute or two, then switch back to the Notes and refresh the views in the SMTP.BOX and Outbound Work Queue. Now, both messages are gone.

The conversion process doesn't concern itself with the message after it's converted. The scope of responsibility for conversion services is very specific. Document transmission and database cleanup are the responsibilities of other MTA tasks.

Exercise #3: Notes-to-Internet message mapping

This final exercise attempts to answer the question, "How do the various fields in a Notes message correspond to those in an Internet message?" The best way to answer this question, is to compare a Notes message to its converted "Internet-ready" counterpart.

1. Choose File - Database - Open and open the Outbound Work Queue database.
2. If the SMTP MTA isn't running on your server, enter LOAD SMTPMTA at the server console.
3. Once the MTA is initialized, enter TELL SMTPMTA STOP DRTASK at the server console. This command stops the Delivery Report Task (DRT task), which is responsible for cleaning up the Outbound Work Queue.
4. Send a message to someone on the Internet and remember to save a copy in your mail database. In your favorite word processor, create a three-column table with one column for values, one for Notes field labels, and one for RFC 822 field labels. To see a sample table that I created during this exercise, see the sidebar "[Notes-to-Internet message mapping](#)."
5. Open the Sent view in your mail database and select the message you just sent.
6. Choose Edit - Properties.
7. Select the Field tab.
8. Add each field shown in this dialog box to the Notes field column of your table, and then add its corresponding value to the value column. If a field doesn't have a value, leave it blank.
9. Switch to the Outbound Work Queue and notice that your message appears under the category WQ_SUCCESS_XMIT, which indicates the message was successfully transmitted.
10. Open this message and notice the Envelope Info section. This is the first part of the three-part Internet message format.
11. Scroll to the bottom of the document and notice the WQBodyHeader0 and WQBody0 attachments.

12. Double-click on the WQBodyHeader0, select View and you'll see the message header created by the conversion task. Analyze this header and look for the values already present in your table. Once you find one, add the field label associated to this value into the RFC 822 field column.
13. Close the header screen and double-click on the WQBody0, select View and you'll see the body section created by the conversion task. Notice the MIME boundary identifiers, which are added when the MTA is configured with a Message Content setting of "Users without Lotus Notes."

Once you've completed this field mapping for an outbound message, try repeating the steps for an inbound message.

Where to next

If the conversion services could speak, they'd probably say, "We just convert messages. We don't deliver them. We don't delete them when they're no longer needed. WE JUST CONVERT MESSAGES!" This single-minded focus accounts for about half of what the MTA does. The other half is message reception and delivery, which is handled by the SMTP MTA transport services ([as we covered in the first article](#)). These two functional areas make up nearly all of the MTA; however, there are still a few other important tasks that help the MTA do its job. These tasks, the Add-In and DRT task, are covered in the final article that puts the SMTP MTA Under the Microscope.

Copyright 1998 Iris Associates, Inc. all rights reserved.