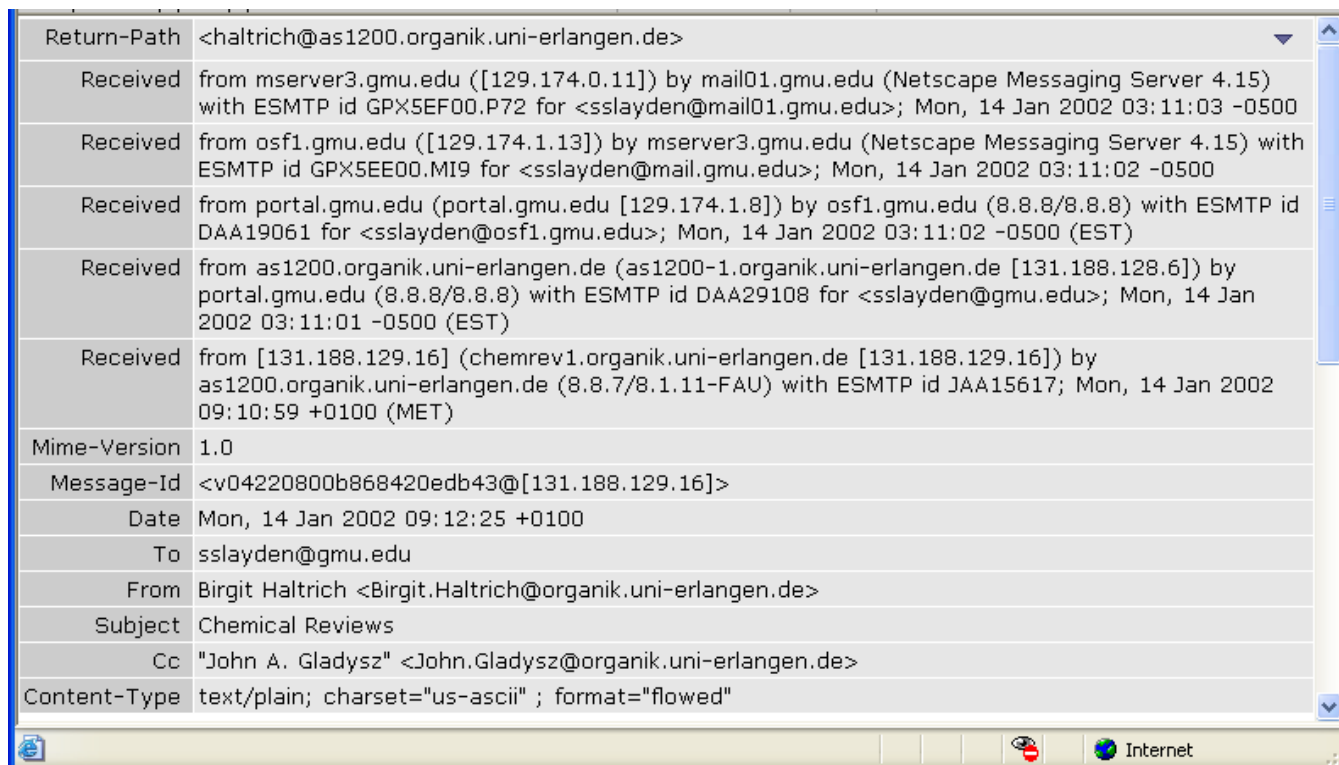


SMTP Moves The Mail: Keeping E-mail Simple

Using electronic mail is such a common occurrence these days that it seems as if everyone has staked his or her place in cyberspace. E-mail actually is the most popular Internet application, and it doesn't appear that anything's going to change that. Sending messages as electrons has become a part of the domestic landscape. Since you're certainly an e-mail user yourself, you should know a little something about the mechanisms behind it.

The basic mechanism for moving e-mail around is called the **Simple Mail Transfer Protocol**, or **SMTP**. Jonathon B. Postel invented the SMTP protocol in 1982, long before the Internet started to become what it is today. At that time the Internet was known as DARPA net, the Defense Advanced Research Projects Agency network. It consisted of about 400 computers in government agencies and universities scattered around the United States. There was already a very basic method for moving text messages, but it had a lot of shortcomings. There was a definite need for a reliable, consistent way of communicating, and SMTP answered that need.

SMTP is just what its first name suggests, simple. A basic e-mail consists of two parts: a "header", which carries a lot of information including the addresses of the sender and receiver; and the "body", which carries the actual message. As a message moves through the Internet, any mail server which handles it adds its own "stamp" to the header, so the recipient can see not only where the message came from, but also what route it took to get there. In Messenger Express, you can see the detailed header in a message by clicking on the little triangle arrow in the top right of the "From:" line of a message. See an example in Figure 1.




Internet Mail Addressing: The Header. Headers can be fairly detailed. The “Return-Path”, shows the destination for any reply to the attached e-mail. This speeds up responses by removing the need for re-addressing the e-mail. In this case, the sender is from Germany (denoted by the ".de" in the return address).

The “**Received**” lines give us a lot of very useful information. They show that from the time the message was first sent in Germany, until it was received by me, it was passed along to many computers, such as *mserver.gmu.edu* and *osfl.gmu.edu*. Each of these machines has an associated Internet Protocol address (IP address). A unique internal ID number, and a time stamp was applied as the message was routed through each machine. Now we know the full route the message took, how long it took to get there, and a fair amount about the sender’s and receiver’s server.

The line, “**Message-Id**” is a unique number that the sending server created to go with the message as it traveled across the Internet. If for some reason the complete message didn’t make it across, the receiving server would ask for it to be resent using that ID.

The next line, “**Date**” gives the time and date the message was placed in the recipient’s “Inbox”. After that is the recipient’s fully qualified (complete) Internet address, which has to include their username (“sslayden”) and the domain they belong to (“gmu.edu”). The “**From**” line includes both the sender’s plain name and their fully qualified e-mail address, and the “**Subject**” line is self-explanatory. Depending on the sender and receiver, there might be many more lines. At GMU, we have a junk mail filter, Ironport, which occupies many lines.

MIME: The original version of SMTP worked great from its invention until the early 1990s. Virtually all information stored on the Internet was text, so moving it from one place to another was no great trick. It didn’t stay that way. In 1989 the very first World Wide Web server and browsers were introduced, and all of a sudden, the Internet became a very colorful, graphical place to be. SMTP needed to be changed so that e-mails could contain not only text, but pictures, sound clips, or other multimedia components. In 1992, Nathaniel Borenstein and Neil Freed submitted the **Multimedia Internet Mail Extensions** (MIME) for approval as an Internet standard. MIME was a very elegant way to handle information of many different types in a single document, and it gained acceptance rapidly. MIME truly is an extension to SMTP. In the e-mail header example, all that is needed to make it capable of transporting multimedia is to add two lines to the header, as in Figure 2.



The first line tells the receiving mail server that all the multimedia commands are compatible with the earliest version of MIME so there won’t be any surprises when the server starts to read in the body of the e-mail. The second line tells the receiving server that when it sees the symbols "boundary=-----090002080806010602030100", then that’s where the multimedia portion begins. The multimedia file itself says how long it is, so the server knows when to “chop off” the stream of data. There can be an unlimited number of multimedia components in a single e-mail, but the real limit is keeping the e-mail to a manageable length.

Now the e-mail is on your server, whether the server is at an Internet Service Provider, your own company, or even your own home. You still need to get to that mail to your computer, so there are three different methods to choose from. The old-fashioned way is to “log on” to the server and use an e-mail program running on the server to view, manage and reply to your mail. Accessing mail directly on the server worked great in the days when minicomputers were all that was available, but was a poor way to use the personal computers that came later. PCs are capable of retrieving, managing and storing e-mails themselves, taking some of the load off the mail server. The only issue is how to get to that mail, so that’s where two other standards come in. The **Post Office Protocol (POP)** has been used as a standard way of sending, checking, retrieving and deleting mail from a centralized server. Most modern e-mail programs use POP, but there’s one problem with it. All mail is stored on local PCs, so there’s no backup and very little security. Enter the **Internet Mail Access Protocol (IMAP)**. IMAP works very similarly to POP, except that e-mails are stored on the server, and the PC e-mail program tells the server how to manage those e-mails rather than managing them itself.

Now you know all about how e-mail gets from here to there, and if things go wrong you can make an educated guess about just where it happened. Just remember one thing: When you send mail, it’s forever. You’ll notice that nowhere in this flow is there a place to say, “Wait! I take it all back!” so be sure you really want to say what you said when you hit that “send” button.