

DISTRIBUTED ELECTRONIC MAIL MODELS IN IMAP4

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Distributed Electronic Mail Models

There are three fundamental models of client/server email: offline, online, and disconnected use. IMAP4 can be used in any one of these three models.

The offline model is the most familiar form of client/server email today, and is used by protocols such as POP-3 (RFC 1225) and UUCP. In this model, a client application periodically connects to a server. It downloads all the pending messages to the client machine and deletes these from the server. Thereafter, all mail processing is local to the client. This model is store-and-forward; it moves mail on demand from an intermediate server (maildrop) to a single destination machine.

The online model is most commonly used with remote filesystem protocols such as NFS. In this model, a client application manipulates mailbox data on a server machine. A connection to the server is maintained throughout the session. No mailbox data are kept on the client; the client retrieves data from the server as is needed. IMAP4 introduces a form of the online model that requires considerably less network bandwidth than a remote filesystem protocol, and provides the opportunity for using the server for CPU or I/O intensive functions such as parsing and searching.

The disconnected use model is a hybrid of the offline and online models, and is used by protocols such as PCMAIL (RFC 1056). In this model, a client user downloads some set of messages from the server, manipulates them offline, then at some later time uploads the changes. The server remains the authoritative repository of the messages. The problems of synchronization (particularly when multiple clients are involved) are handled through the means of unique identifiers for each message.

Each of these models have their own strengths and weaknesses:

Feature	Offline	Online	Disc
-----	-----	-----	----
Can use multiple clients	NO	YES	YES
Minimum use of server connect time	YES	NO	YES
Minimum use of server resources	YES	NO	NO
Minimum use of client disk resources	NO	YES	NO
Multiple remote mailboxes	NO	YES	YES
Fast startup	NO	YES	NO
Mail processing when not online	YES	NO	YES

Although IMAP4 has its origins as a protocol designed to accommodate the online model, it can support the other two models as well. This makes possible the creation of clients that can be used in any of the three models. For example, a user may wish to switch between the online and disconnected models on a regular basis (e.g. owing to travel).

IMAP4 is designed to transmit message data on demand, and to provide the facilities necessary for a client to decide what data it needs at any particular time. There is generally no need to do a wholesale transfer of an entire mailbox or even of the complete text of a message. This makes a difference in situations where the mailbox is large, or when the link to the server is slow.

More specifically, IMAP4 supports server-based RFC 822 and MIME processing. With this information, it is possible for a client to determine in advance whether it wishes to retrieve a particular message or part of a message. For example, a user connected to an IMAP4 server via a dialup link can determine that a message has a 2000 byte text segment and a 40 megabyte video segment, and elect to fetch only the text segment.

In IMAP4, the client/server relationship lasts only for the duration of the TCP connection. There is no registration of clients. Except for any unique identifiers used in disconnected use operation, the client initially has no knowledge of mailbox state and learns it from the IMAP4 server when a mailbox is selected. This initial transfer is minimal; the client requests additional state data as it needs.

As noted above, the choice for the location of mailbox data depends upon the model chosen. The location of message state (e.g. whether or not a message has been read or answered) is also determined by the model, and is not necessarily the same as the location of the mailbox data. For example, in the online model message state can be co-located with mailbox data; it can also be located elsewhere (on the client or on a third agent) using unique identifiers to achieve

common reference across sessions. The latter is particularly useful with a server that exports public data such as netnews and does not maintain per-user state.

The IMAP4 protocol provides the generality to implement these different models. This is done by means of server and (especially) client configuration, and not by requiring changes to the protocol or the implementation of the protocol.

Security Considerations

Security issues are not discussed in this memo.

Author's Address:

Mark R. Crispin
Networks and Distributed Computing, JE-30
University of Washington
Seattle, WA 98195

Phone: (206) 543-5762

EMail: MRC@CAC.Washington.EDU