

EMSD: The LEAP E-Mail Component

Mohsen Banan

<http://mohsen.banan.1.byname.net/ContactMe>

Version 1.1

First Published: August 4, 2000

Last Updated: July 14, 2000

Copyright ©2000 Mohsen Banan

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

A Component of *The LEAP Manifesto*

This article is one of a series of articles describing various aspects of the Mobile Messaging industry and the Lightweight & Efficient Application Protocols (LEAP) protocols. For the complete collection of articles see *The LEAP Manifesto* [2], available at

<http://www.LeanForum.org/LEAP/Manifesto/roadMap/index.html>. *The LEAP Manifesto* is also available at the Free Protocols Foundation website at

<http://www.FreeProtocols.org/LEAP/Manifesto/roadMap/index.html>.

Contents

1	Introduction	3
1.1	Terminology	3
2	Existing Internet Mail Submission and Delivery	3
3	Overview of EMSD	4
3.1	Protocol Layering	4
3.2	EMSD Protocol Components	4
3.3	Efficient Short Remote Operations (ESRO)	5
3.4	Anticipated Uses of EMSD	5
4	EMSD Design Goals and Requirements	6
5	Rationale for Key Design Decisions	7
5.1	Deviation from the SMTP Model	7
5.1.1	Efficiency Comparison of SMTP and EMSD	7
5.2	Use of ESRO Instead of TCP	7
5.3	Use of the Remote Procedure Call (RPC) Model	8
5.4	Use of ASN.1	8
6	Relationship of EMSD to Other Mail Protocols	9
7	Obtaining the EMSD Protocols	9

List of Figures

1	LEAP Protocol Stack	11
2	Efficient Mail Submission and Delivery Protocol	12
3	EMSD World and Global Messaging World	13
4	Messaging Communication Stack and EMSD	14

List of Tables

1	Comparison of EMSD to Other Protocols	8
2	Messaging Protocol Functionality	9

1 Introduction

This article provides a description of the **Efficient Mail Submission & Delivery** protocol, or **EMSD**. EMSD is the e-mail component of the LEAP family of protocols.

The entire family of LEAP protocols has been designed with efficiency as a primary requirement, and each component protocol brings efficiency and functionality benefits to the users of miniaturized mobile devices. In particular, EMSD brings these efficiency benefits to the e-mail, or Mobile Messaging, application.

Mobile Messaging is just one of several applications for which there is a high efficiency premium in the mobile and wireless arena. Other applications which demand efficiency in the mobile environment are such things as web browsing, dictionary look-up, etc. From an industry-building perspective, however, e-mail is the most critically important of all wireless and mobile data communications applications. Mobile Messaging provides the user with a unique communications facility: the immediate delivery of important and time-critical information to the mobile recipient, wherever and whenever he/she happens to be. This is a capability which is not provided by phone, Fax, or any other tool in the modern user's array of communications options.

In fact, this capability represents the principal value of wide-area wireless networks to the end user, and this is why e-mail remains the dominant application for wide-area wireless networks. Furthermore, miniature hand-held mobile devices are extremely well-suited to the e-mail application. The same is not true for other applications, such as web browsing.

For all these reasons, e-mail is the key industry-building application, and this is why EMSD represents the right starting point for the LEAP family of protocols.

1.1 Terminology

Throughout this article, we will make use of the following terms and definitions:

- **MTS**. Mail Transfer System.
- **Mail Submission**. Mail submission refers to the process of putting mail into the Mail Transfer System (MTS).
- **Mail Delivery**. Mail delivery refers to the process of the MTS putting mail into a user's final mailbox.
- **SMTP**. Simple Mail Transfer Protocol.
- **MUA**. Mail User Agent.
- **EMSD-P**. EMSD Protocol.
- **EMSD-FS**. EMSD Format Standard.
- **EMSD-UA**. EMSD User Agent.
- **EMSD-SA**. EMSD Server Agent.

2 Existing Internet Mail Submission and Delivery

Mail transmission in the Internet did not arise as a result of well-planned engineering processes; rather, it grew and evolved in a more organic way.

At present, most mail submission and delivery throughout the Internet is done by means of the Simple Mail Transfer Protocol, or SMTP. SMTP was originally defined as a message *transfer* protocol – that is, a means to route

(if necessary) and deliver mail by putting finished (i.e. complete) messages in a mailbox. Originally, users connected to servers from terminals, and all processing occurred on the server. Today, a split-MUA (Mail User Agent) model is common, in which MUA functionality occurs both on the user's own system, and the server.

In the split-MUA model, the process of getting a message to the user is accomplished by access to a mailbox on the server, using protocols such as POP and IMAP. Also, in the split-MUA model, the user's access to his/her message is based on a "Message Pull" paradigm, in which the user is required to explicitly poll his/her mailbox to retrieve mail. Message delivery based on a "Message Push" paradigm, in which mail is delivered directly to the user without polling, is presently not supported.

Despite its original definition as a message *transfer* protocol, in the split-MUA model, SMTP is often used for message *submission*. The widespread use of SMTP for submission has become a reality, regardless of whether this is a good or a bad thing.

3 Overview of EMSD

EMSD is a messaging protocol that is highly optimized for the submission and delivery of short Internet e-mail messages. The EMSD protocol addresses all the shortcomings in the existing Internet mail system described in the previous section. EMSD properly supports the Message Push mode of operation, and it provides an alternative mechanism to SMTP for message submission. And most important of all, it does this with a major emphasis on efficiency.

3.1 Protocol Layering

As shown in Figure 1, the LEAP protocols are layered. The lower layer, called Efficient Short Remote Operations (ESRO), provides efficient reliable connectionless transport services which can be used by a variety of applications. For example, in addition to Mobile Messaging services, ESRO can also be used as a transport service for credit card verification applications and efficient micro browsers.

EMSD is built on top of ESRO. The reliability requirements for message submission and message delivery in EMSD are the same as for existing e-mail protocols. The EMSD protocol provides reliable connectionless mail submission and delivery services.

3.2 EMSD Protocol Components

EMSD consists of two independent components: the **EMSD Format Standard**, and the **EMSD Protocol**. These two components provide the following functions:

1. **EMSD Format Standard (EMSD-FS)**

EMSD-FS is a non-textual form of compact encoding of Internet e-mail (RFC-822) messages, which facilitates efficient message transfer. EMSD-FS is used in conjunction with the EMSD-P (described below), but is not in any way a general replacement for RFC-822. EMSD-FS defines a method of representation of short interpersonal messages. It defines the "Content" encoding (Header + Body). Although EMSD-FS contains end-to-end information, its scope is purely point-to-point. EMSD-FS relies on EMSD-P for the transfer of the content to its recipients.

2. **EMSD Protocol (EMSD-P)**

EMSD-P is responsible for wrapping a limited size EMSD-FS message in a point-to-point envelope, and submitting or delivering it. EMSD-P performs the envelope encoding. EMSD-P relies on the services of Efficient Short Remote Operations (ESRO) as specified in RFC-2188 [1] for transporting the point-to-point envelope.

Some of the services provided by EMSD-P include: message originator authentication, and optional message segmentation and re-assembly. EMSD-P is expressed in terms of abstract services using the ESRO notation.

Together, the EMSD Protocol and Format Standard define the protocols used to transfer messages between an EMSD Server Agent (EMSD-SA), for example a Message Center, and an EMSD User Agent (EMSD-UA), for example a Two-Way Pager.

Figure 2 illustrates how EMSD defines the communication between a specific EMSD-UA and a specific EMSD-SA. The Message Transfer System may include a number of EMSD-SAs, and each EMSD-SA may have any number of EMSD-UAs with which it communicates.

It is important to note that EMSD-P and EMSD-FS are not end-to-end, but instead focus on the point-to-point transfer of messages. The two points being referred to are the EMSD-SA and EMSD-UA. EMSD-P functions as an element of the Internet mail environment, providing end-to-end services (i.e. EMSD-User to any other Messaging Originator or Recipient).

The EMSD services use the Efficient Short Remote Operations (ESRO) services. They also use the Duplicate Operation Detection Support Functions. These functions guarantee that an operation is performed no more than once.

3.3 Efficient Short Remote Operations (ESRO)

The EMSD protocol specifications define the protocols between the EMSD Device and the EMSD Server. EMSD is built on top of, and requires the services of, ESRO (Efficient Short Remote Operations). This EMSD requirement was the major motivation for the development of ESRO; however, ESRO has been developed to be independent of EMSD.

ESRO defines a notation and the services provided by an application-service element to support interactive applications in a distributed systems environment. The scope of ESRO services is not limited to EMSD. ESRO is designed to be able to support other applications, such as finger/limited directory service.

The ESRO protocol provides reliable connectionless remote operation services on top of UDP (or any other non-reliable connectionless transport service) with minimum overhead. ESRO supports segmentation and reassembly, concatenation and separation, as well as multiplexing for service users (i.e. applications).

The ESRO service is similar to and is consistent with other Remote Procedure Call services. The major emphasis of the ESRO service definition and the ESRO protocol is on efficiency. ESRO has been designed specifically with wireless network (e.g. CDPD) usage in mind.

The service model, the notation, and the protocol for ESRO are fully specified in RFC-2188 [1]. The EMSD Protocol uses ESRO to accomplish reliable connectionless mail submission and delivery.

For more information on ESRO, see the article entitled *ESRO: A Foundation for the Development of Efficient Protocols* within *The LEAP Manifesto*, or visit the ESRO website at <http://www.esro.org/>.

3.4 Anticipated Uses of EMSD

Any network or network operator which faces significant bandwidth and capacity limitations can benefit from the use of EMSD. Any user of a network who must bear high costs for measured network usage can benefit from the use of EMSD.

The initial use of EMSD is expected to be primarily to provide Mobile Messaging services over IP-based wireless networks. However, EMSD can also function as an adjunct to Mail Access Protocols for "Mail Notification Services."

Mail submission and delivery take place at the edges of the network. It is likely that multiple mail submission and delivery protocols will be developed, each addressing the specific requirements of a particular user's environment. Such diversity on the edges of the network is beneficial, and with the right protocols, this diversity does not adversely

affect the integrity of the mail transfer system. EMSD is the basis for the mail submission and delivery protocol to be used when the user's environment demands efficiency.

4 EMSD Design Goals and Requirements

The EMSD protocols have been designed to accomplish three high-level goals:

1. Define the new "world" of Efficient Mail Submission & Delivery
2. Define a remote operations service that can handle messaging and other standard networking applications
3. Make EMSD an extension of the existing Internetworking world

Based on these goals, EMSD has been designed to satisfy the following design requirements:

1. Support the submission of short mail messages with the same (or better) level of functionality as the existing Internet mail protocols.
2. Support the delivery of short mail messages with the same (or better) level of functionality as the existing Internet mail protocols.
3. Function as an extension of the existing mainstream Internet mail.
4. Minimize the number of transmissions.
5. Minimize the number of bytes transmitted.
6. Be quick: minimize the latency of message submission and delivery.
7. Provide the same level of reliability (or better) as the existing e-mail protocols.
8. Accommodate varying sizes of messages: the size of a message may determine how the system deals with the message, but the system must accommodate it.
9. Be power efficient and show respect for mobile platform resources, including memory and CPU levels, as well as battery power longevity. In other words, be client-light and server-heavy.
10. Be highly extensible. Different users will demand different options, so the solution cannot require every feature to be a part of every message. Likewise, usage will emerge that is not currently recognized as a requirement. The solution must be extensible enough to handle new, emerging requirements.
11. Be secure. Provide the same level of security (or better) as the existing e-mail protocols. Content confidentiality, originator/recipient authentication, and message integrity must be available options to users.
12. Be easy to implement: re-use existing technology as much as possible.

The EMSD protocols make extensive use of existing technology, including:

- RFC-822
- ASN.1
- Basic Encoding Rules

- Internet mail

By using these established technologies, the design of EMSD avoids the expense and other problems associated with "re-inventing the wheel." The above technologies have been thoroughly tested, and have proven to be reliable solutions for the problems they address (e.g. message format, reliable message delivery, encoding and compacting). The EMSD specifications cater to users who enjoy the advantages of this new technology, but at the same time want to be connected to the rest of the existing Internet e-mail world. Figure 3 shows how the Global and EMSD worlds complement one another.

The Internet e-mail community is shown in the lower half of the figure. This world is connected to the EMSD Internet e-mail system.

5 Rationale for Key Design Decisions

This section summarizes the rationale for the key design decisions that were made while developing the EMSD protocols.

5.1 Deviation from the SMTP Model

SMTP is the main mail transport mechanism used throughout the Internet. It is widely deployed and well understood by many engineers who specialize in Internet e-mail. For these reasons, protocols based on or derived from SMTP or more likely to become widely deployed throughout the Internet.

However, SMTP is highly inefficient for the transfer of short messages. SMTP is inefficient both in terms of the number of transmissions, and in terms of the number of bytes transmitted. Even when fully optimized with PIPELINING, SMTP remains significantly inefficient.

The submission of a short message using SMTP requires 15 transmissions. The submission of a short message with SMTP and PIPELINING requires 9 transmissions. The submission of a short message with EMSD (EMSD-P and ESRO) requires only 3 transmissions (in a typical case).

The key design requirement of EMSD is efficiency. Because of the 3 fold (at least) gain in efficiency, this justifies the deviation from the SMTP model.

5.1.1 Efficiency Comparison of SMTP and EMSD

Table 1 shows the number of N-PDUs exchanged for the transfer of a short Internet e-mail when using SMTP, SMTP with PIPELINING, QMTP, and EMSD. The names used for identifying the PDUs are informal names.

5.2 Use of ESRO Instead of TCP

In order to provide the same level of reliability that the existing e-mail protocols provide for short messages, it is clear that a reliable underlying service is needed. UDP [3], by itself, is clearly not adequate.

Use of TCP however, involves three phases:

1. Connection Establishment
2. Data Transfer
3. Disconnect

	SMTP	SMTP + Pipelining	QMTP , QMQP	EMSD
Client:	SYN	SYN	SYN	Submit Req
Server:	SYN ok	SYN ok	SYN	Submit Resp
Client:	HELLO	HELLO	message	ack
Server:	ok	PIPELINING	accept close	
Client:	MAIL	MAIL RCPT DATA	close	
Server:	ok	ok		
Client:	RCPT	message QUIT		
Server:	ok	accept ok close		
Client:	DATA	close		
Server:	ok			
Client:	message			
Server:	accept			
Client:	QUIT			
Server:	ok close			
Client:	close			

Table 1: Comparison of EMSD to Other Protocols

The reliable transfer of a short message using TCP involves a minimum of five transmissions, as is the case with QMTP.

Again, the key design requirement of EMSD is efficiency. Therefore deviation from TCP is justified, because this eliminates the two extra transmissions that are an inherent characteristic of TCP.

The ESRO protocol, as specified in RFC-2188 [1], provides reliable connectionless remote operation services on top of UDP [3] with minimum overhead. ESRO supports segmentation and reassembly, concatenation and separation.

The reliable transfer of a short message using ESRO involves 3 transmissions, as is the case with EMSD-P.

5.3 Use of the Remote Procedure Call (RPC) Model

Many Internet protocols are "text-based." On the other hand, few Internet protocols are RPC-based. Protocols designed on the basis of the "text-based" approach have a better track record of acceptance throughout the Internet.

However, considering that message submission and delivery in EMSD involves no more than two data exchanges, the text-based model becomes the same as an operation. Furthermore, the RPC model is the natural way of using ESRO.

5.4 Use of ASN.1

In order to minimize the number of bytes transferred, efficient encoding mechanisms are needed.

Among today's encoding mechanisms, ASN.1 has the unique feature of separating the abstract syntax from the encoding rules. By selecting ASN.1 as the notation used for expressing EMSD's information objects, EMSD has the flexibility of using the most efficient encoding rules, such as Packed Encoding Rules (PER), when they are available.

Efficient encoding can always be better performed when the syntax of the information is known. In general, encoding and compression techniques which use the knowledge of the syntax of the information produce better results than those compression techniques that work on arbitrary text.

Protocol Function	SMTP	IMAP	POP	EMSD
Submission	XX			XXX
Delivery	XXX			XXX
Relay (Routing)	XXX			
Retrieval		XXX	XXX	XX
Mailbox Access		XXX	X	
Mailbox Sync.		XXX		

Table 2: Messaging Protocol Functionality

6 Relationship of EMSD to Other Mail Protocols

EMSD is designed to be a companion to existing Internet mail protocols. It is designed to fit within the many protocols already in use for messaging, as well as those already in use for networking.

Figure 4 shows how EMSD fits in with the other major messaging protocols. The RFCs referenced in the figure are current at the time of this writing, but could be updated or made obsolete at any time.

The various Internet mail protocols provide different sets of capabilities for mail processing. Table 2 summarizes the capabilities of SMTP, IMAP, POP and EMSD in the following areas of functionality:

- Mail Submission
- Mail Delivery
- Mail Routing (Relay)
- Mail Retrieval
- Mailbox Access
- Mailbox Synchronization

The number of X's in each cell of the table denotes the extent to which a particular function is supported by a particular protocol.

Table 2 clearly demonstrates that combinations of these protocols can be used to complement one other in providing rich functionality to the user. For example, a user interested in highly mobile messaging functionality can use EMSD for the submission and delivery of time-critical and important messages, and use IMAP for comprehensive access to his/her mail-box.

7 Obtaining the EMSD Protocols

For complete instructions on how to obtain the EMSD protocols, visit the "Base Protocol Specifications" section of <http://www.emsd.org/>.

References

- [1] M. Banan, M. Taylor, and J. Cheng. AT&T/Neda's Efficient Short Remote Operations (ESRO) Protocol Specification Version 1.2. RFC 2188 (Informational), September 1997.

- [2] Mohsen Banan. *Lightweight & Efficient Application Protocol (LEAP) Manifesto*. Technical Report 108-101-01, LEAP Forum, Bellevue, WA, January 2000. Online document is available at <http://www.leapforum.org/LEAP/Manifesto/completeManifesto>.
- [3] Jon B. Postel. User Datagram Protocol. Request for Comments 768, DDN Network Information Center, SRI International, August 1980.

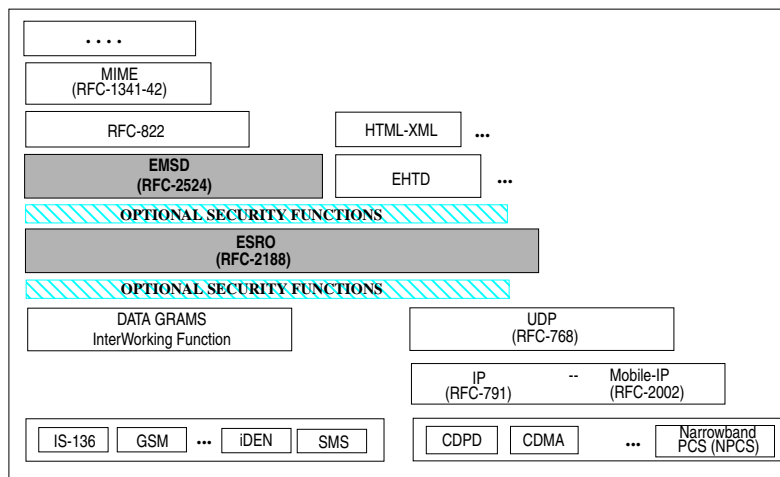


Figure 1: LEAP Protocol Stack




Figure 2: Efficient Mail Submission and Delivery Protocol

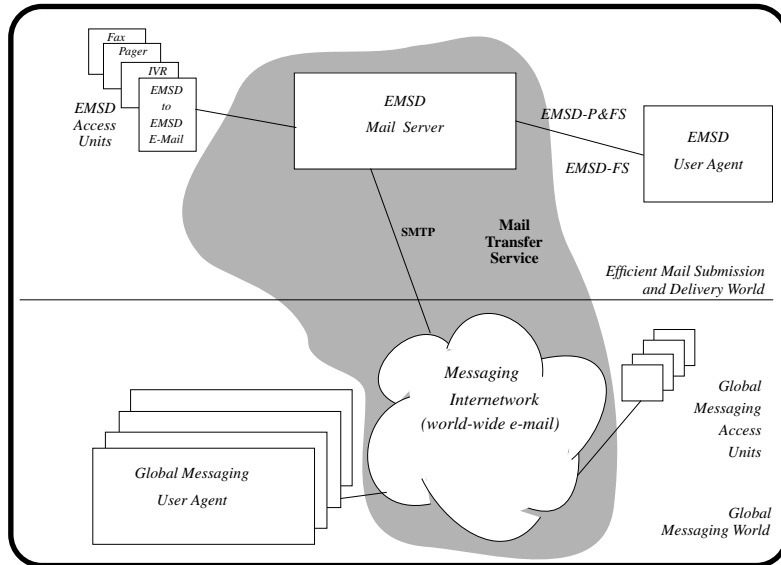


Figure 3: EMSD World and Global Messaging World

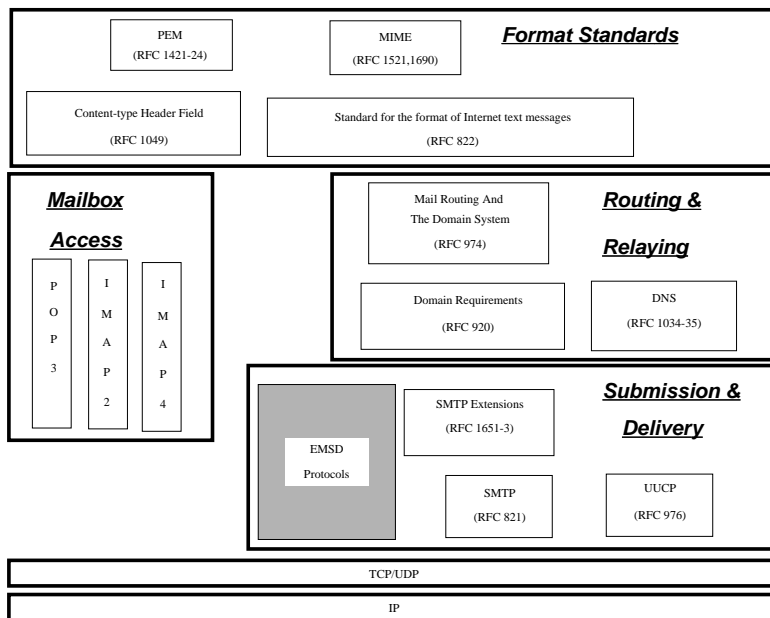


Figure 4: Messaging Communication Stack and EMSD