WHITEPAPER SERIES

Post Office Protocol Version 3 (POP3) and Internet Electronic Mail

Version 1.0 **November 1998**

Hong Kong Computer Center, 20/F
54-62 Lockhart Road
Wan Chai
HONG KONG
Tel: +852 2520-0300
Fax: +852 2648-5913

The Peak Tower, 15/F
107 Alfaro Street
Salcedo Village, Makati City
PHILIPPINES
+63 (2) 811-3999
+63 (2) 811-3939

USA Support/Sales: +1 (408) 481-9985
USA Fax: +1 (888) 562-3561

Email: info@ima.com
Website: http://www.ima.com

Please send all comments and suggestions to the Whitepaper Series to *doc@ima.com*

TABLE OF CONTENTS

Copyright © 1998 IMA Philippines, Inc. All rights reserved

INTRODUCTION

The Post Office Protocol version 3 (POP3), as defined in RFC 1939, is an "offline" protocol designed to provide the Internet community with a tool for connecting to mail servers and retrieving email messages. Since the Simple Mail Transfer Protocol (SMTP) functions only as a transport protocol between mail servers, it is not capable of delivering mail from the post office to a user's mailbox. It also does not allow remote users to get mail from the server. Thus, a delivery protocol like POP3 is needed so that the mail received by the server from the Internet can be delivered to their intended recipients.

As an offline protocol, POP3 retrieves messages from the server (either at an ISP or on a LAN) and transfers them to a workstation's hard disk. It deletes the messages from the server if this option is explicitly mentioned in the configuration. In short, POP3's function is to get email from a remote mailbox and store it on a user's local machine so it can be read later in a disconnected or "offline" state. POP3 is designed as a single user, single mailbox system (one account per user). The POP3 connector is usually a combined POP3 retriever and SMTP sender.

With POP3, smaller nodes in the Internet that are incapable of maintaining a message transport system (MTS) can retrieve mail. An example of such nodes are workstations that do not have the resources needed to maintain an SMTP server and associated mail delivery system. Another example is a personal computer that connects to the Internet only at certain times due to economic reasons. A node capable of supporting an MTS can offer maildrop service to these smaller nodes via POP3 service. Through this setup, a workstation can dynamically access a maildrop on a host that offers the POP3 service, otherwise known as the server host. The host that makes use of the POP3 service is called the client host.

The server host starts the POP3 service by listening to TCP port 110. Users who want to retrieve their mail must log in to the server host POP3 service at this port with their account names and passwords. The client host wishing to use the POP3 service establishes a TCP connection with the server host. After the connection has been made, the POP3 server sends a greeting. The client and the POP3 server then issue commands and responses until the connection is terminated.

All POP3 commands are terminated by a CRLF pair. These commands consist of a keyword, possibly followed by one or more arguments. Each of the keywords and arguments, which consist of printable ASCII characters, is separated by a single SPACE character. Keywords may be three or four characters long, while an argument may have a maximum length of 40 characters. Responses are also terminated by a CRLF pair. Responses in POP3 consist of a status indicator and a keyword, which may be followed by additional information. There are currently two status indicators – positive ("OK") and negative ("ERR"). The status indicators are case-sensitive and must be sent in upper case by the server.

THE POP3 SESSION

The POP3 session goes through a number of states during its lifetime (see Figure 1). After TCP connection has been established and the POP3 server has sent a greeting, the session enters the AUTHORIZATION state. While in this state, the client identifies itself to the POP3 server. After this stage, the server acquires resources associated with the client's maildrop. The session then enters the TRANSACTION state, wherein the

Copyright © 1998 IMA Philippines, Inc. All rights reserved

client requests actions on the part of the server. Once the client has issued the QUIT command, the POP3 server releases all the resources acquired during the TRANSACTION state and says goodbye. The TCP connection is then terminated.

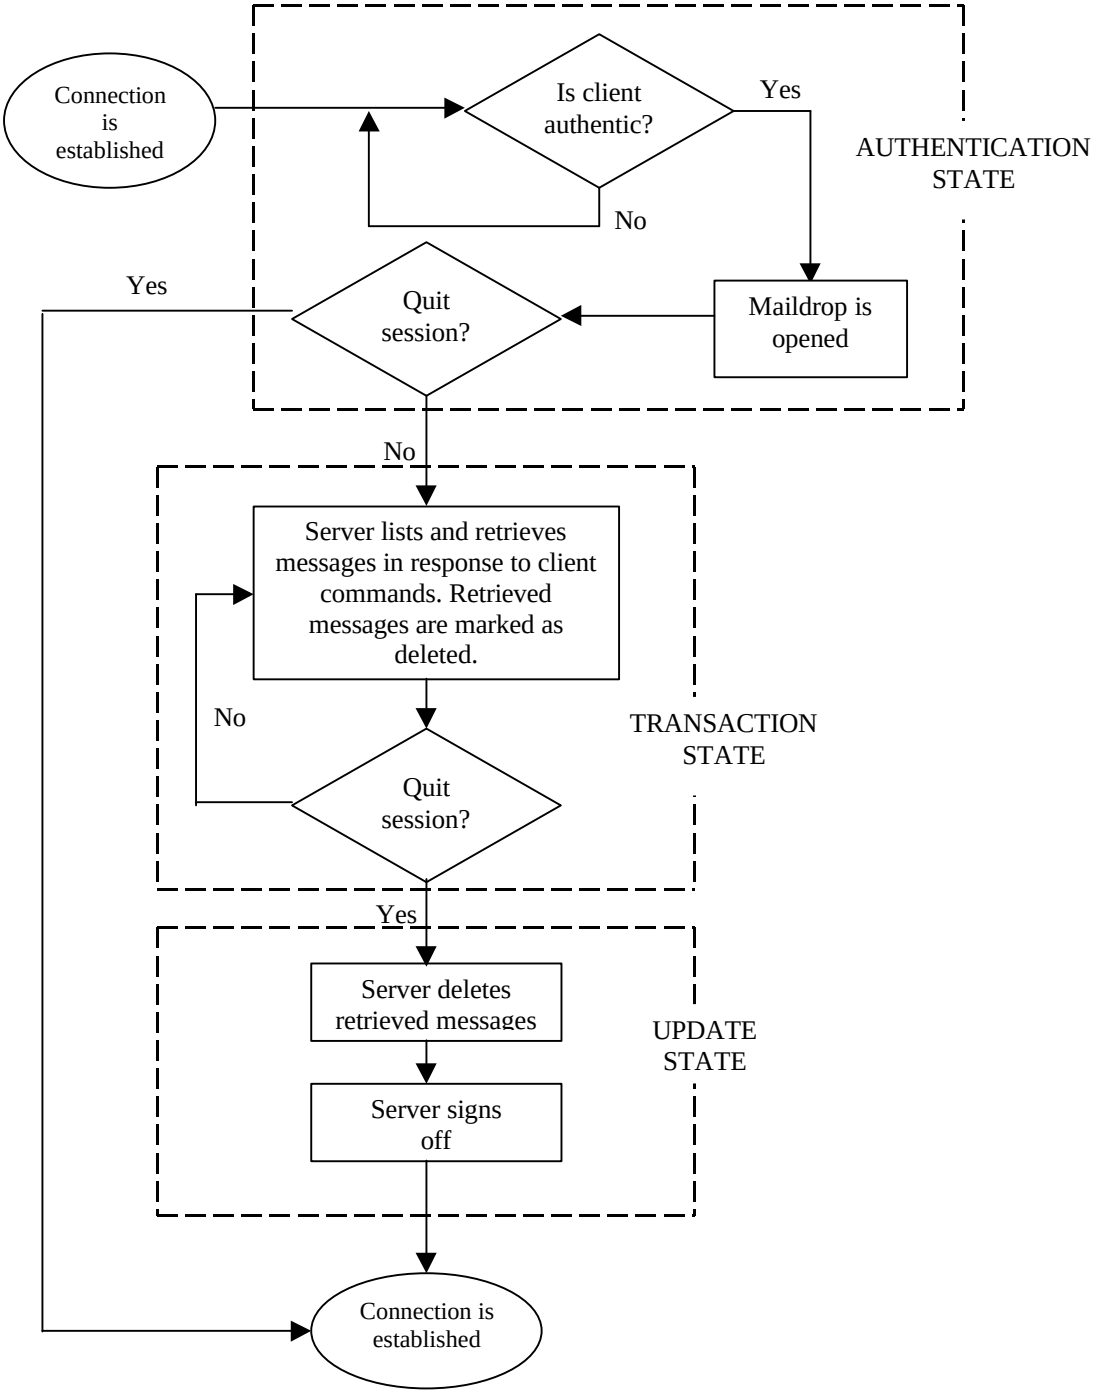

Figure 1

When a client gives an unrecognized, unimplemented or syntactically invalid command, the server automatically responds with a negative status indicator. The same response is given when a command is issued in an incorrect state. There is no general method for

Copyright © 1998 IMA Philippines, Inc. All rights reserved

a client to distinguish between a server that does not implement an optional command and another server that is unable or unwilling to process the command.

A POP3 server may have an inactivity autologout timer. Such a timer must be set to a duration of at least 10 minutes. If the server receives a command from the client during that interval, the autologout timer is reset. When the timer expires, the session is prevented from entering the UPDATE state, that is the server terminates the TCP connection without removing any messages or sending any response to the client.

The AUTHORIZATION State

The POP3 server issues a one-line greeting after the TCP connection has been established by a POP3 client. This can be any positive response. An example might be:

S: +OK POP3 server ready

Once the server has issued this one-line greeting, the POP3 session enters the AUTHORIZATION state. In this state, the client is required to identify and authenticate itself to the POP3 server. Two possible mechanisms for doing this are the USER and PASS command combination, and the APOP (Authenticated POP) command. RFC 1734 describes another authentication mechanism, the AUTH command, which is based on the protection mechanisms for IMAP4. A POP3 server must be capable of supporting at least one of these mechanisms.

Once the authenticity of the client has been validated, it is given access to the appropriate maildrop, in which case the POP3 server acquires an exclusive-access lock on the maildrop to prevent the messages from being modified or removed before the session enters the UPDATE state. If the lock is acquired successfully, the POP3 server responds with a positive status indicator. The session then enters the TRANSACTION state, with no messages marked as deleted. If the maildrop cannot be opened for any reasons (for example, a maildrop cannot be parsed or a lock cannot be acquired), the POP3 server responds with a negative status indicator. After receiving a negative response, the client may either issue the QUIT command, or issue a new authentication command and start again (if connection is not terminated by the POP3 server).

After the maildrop has been successfully opened by the POP3 server, it assigns a message number to each message, and notes the size of every message in octets. All message numbers and message sizes are expressed in base-10.

The TRANSACTION State

After the client has been successfully identified and the POP3 server has locked and opened the appropriate maildrop, the POP3 session enters the TRANSACTION state. In this state, the client may issue any of the valid POP3 commands, namely:

STAT – when the client issues this command, the POP3 server issues a positive response with a line that contains information for the maildrop. This line is called a “drop listing” for that maildrop. The line contains the positive response “+OK” followed by a single space, the number of messages in the maildrop, and the size of the maildrop in octets.

Copyright © 1998 IMA Philippines, Inc. All rights reserved

LIST [msg] (an optional message-number referring to a message not marked as deleted) – when the server receives this command, it could either give a negative or a positive response. A positive response will contain a line that serves as the “scan listing” for a particular message. A scan listing contains the message-number of the message, followed by a single space and the exact size of the message in octets.

RETR msg (a required message-number referring to a message not marked as deleted) – if the POP3 server issues a positive response after receiving this command, it will then send the message that corresponds to the given message number, being careful to byte-stuff the termination character.

DELE msg (a required message-number referring to a message not marked as deleted) – after receiving this command, the POP3 server marks the message as deleted. However, the server does not actually delete the message until the POP3 session enters the UPDATE state.

NOOP – the POP3 server does nothing except to reply with a positive response.

RSET – after receiving this command, the POP3 server issues a positive response and unmarks all messages that have been marked as deleted.

These commands may only be given in the TRANSACTION state. After receiving each command, the POP3 server issues a response. An example of a POP3 session in the TRANSACTION state is:

C: STAT
S: +OK 2 320
C: LIST
S: +OK 2 messages (320 octets)
S: 1 120
S: 2 200
S: .

The UPDATE State

Once the client has issued the QUIT command in the TRANSACTION state, the POP3 session enters the UPDATE state. If the client issues the QUIT command from the AUTHORIZATION state, the POP3 session terminates without entering the UPDATE state.

If a session is terminated for some other reason other than a client-issued QUIT command, the POP3 session does not enter the UPDATE state and no messages will be removed from the maildrop. The following is an example of a POP3 session in the UPDATE state:

C: QUIT
S: +OK davao POP3 server signing off (2 messages left)

OPTIONAL POP3 COMMANDS

All minimal implementations of POP3 servers are required support all the POP3 commands previously mentioned. There are also several optional POP3 commands that

Copyright © 1998 IMA Philippines, Inc. All rights reserved

will provide a POP3 client with more flexibility in handling messages while preserving a simple POP3 server implementation. The commands are:

TOP msg n (a required message-number referring to a message not marked as deleted, and a non-negative number of lines) – after the server issues a positive response to this command, it will then send the headers of the message, the blank line separating the headers from the body, and the number of lines of the message's body.

UIDL [msg] (an optional message-number referring to a message not marked as deleted) – after issuing a positive response, the server then gives a line called the "unique-id listing" for the message. The unique-id is an arbitrary server-determined string that consists of one to 70 characters in the range 0x21 to 0x7E. It uniquely identifies the message within a maildrop.

USER name (a required string that identifies a mailbox) – the client must first issue this command for authentication using the USER and PASS command combination. If the server responds with a positive status indicator, the client may then issue either the PASS command to complete authentication or the QUIT command to end the POP3 session.

PASS string (a server/mailbox-specific password) – after receiving this command, the POP3 server determines whether to give access to the client using the argument pair from the USER and PASS commands.

APOP name digest (a string that identifies a mailbox and a MD5 digest string) – an optional command that provides password security to POP3 clients.

These commands may be used only in the TRANSACTION state. As with basic POP3 commands, the server issues either a positive status indicator or a negative status indicator in response to each of the commands. Following are several examples of POP3 sessions using optional POP3 commands:

Example 1

C: TOP 1 10
S: +OK (the POP3 server then sends the headers of the message, a blank line, and the first 10 lines of the message body)
S: .

Example 2

C: USER tkehres
S: +OK tkehres is a real hoppy frood
C: PASS secret
S: -ERR invalid password

The following is an example of a full-length POP3 session using basic and optional POP3 commands:

S: <wait for connection on TCP port 110>

Copyright © 1998 IMA Philippines, Inc. All rights reserved

C: <open connection>
S: +OK POP3 server ready <1896.697170952@ima.com>
C: APOP tkehres c4c9334bac560ecc979e58001b3e22fb
S: +OK tkehres' maildrop has 3 messages (420 octets)
C: STAT
S: +OK 2 320
C: LIST
S: +OK 3 messages (420 octets)
S: 1 120
S: 2 200
S: 3 100
S: .
C: RETR 1
S: +OK 120 octets
S: <POP3 server sends message 1>
S: .
C: DELE 1
S: +OK message 1 deleted
C: RETR 2
S: +OK 200 octets
S: <POP3 server sends message 2>
S: .
C: DELE 2
S: +OK message 2 deleted
C: RETR 3
S: +OK 100 octets
S: <POP3 sends message 3>
S: .
C: DELE 3
S: +OK message 3 deleted
C: RSET
S: +OK tkehres' maildrop has 3 messages (420 octets)
C: QUIT
S: +OK davao POP3 server signing off (maildrop empty)
C: <close connection>
S: <wait for next connection>

SECURITY ISSUES

Normally, a POP3 session begins with a USER/PASS exchange. This results in a server/user-id specific password being sent in the clear over the network. For POP3 client implementations that establish connections to the POP3 server on a regular basis, this poses a serious security hazard since the risk of password capture is significantly increased.

The APOP command allows origin authentication and replay protection without requiring the user of the POP3 client to send a password in the clear over the network. This helps reduce the risk of password capture considerably.

A POP3 server capable of supporting the APOP command includes a timestamp in its banner greeting. This timestamp is changed each time the server issues a new banner

Copyright © 1998 IMA Philippines, Inc. All rights reserved

greeting. The syntax of the timestamp corresponds to the “msg-id” described in RFC 822. For example, on a UNIX implementation where a separate UNIX process is utilized for each instance of a POP3 server, the syntax of the timestamp might appear as:

<process-ID.clock@hostname>

where “process-ID” is the decimal value of the process’ PID, clock is the decimal value of the system clock, and hostname is the fully-qualified domain name (FQDN) that corresponds to the host running the POP3 server.

The POP3 client takes note of this timestamp and then issues an APOP command. The “name” parameter in the APOP name digest command has semantics identical to the “name” parameter of the USER command. To calculate for the digest parameter, the MD5 algorithm is applied to a string consisting of the timestamp (including angle-brackets) followed by a shared secret. This shared secret is known only to the POP3 client and server. Any entity that becomes privy to the shared secret will be able to masquerade as the named user, so it is important that great care be taken to prevent unauthorized disclosure of the secret. The digest parameter itself is a 16-octet value sent in hexadecimal format using lower-case ASCII characters.

After the POP3 server has received the APOP command, it verifies the authenticity of the digest provided. If the digest is correct, a positive response is issued by the server and the POP3 session enters the TRANSACTION state. If the digest provided is found to be incorrect, a negative response is issues and the session remains in the AUTHORIZATION state.

It must be noted that as the length of the shared secret increases, so does the difficulty of deriving it. Thus, the use of long strings as shared secrets is adviced. The following is an example of a POP3 session that makes use of the APOP command:

S: +OK POP3 server ready <1896.697170952@ima.com>
C: APOP tkehres c4c9334bac560ecc979e58001b3e22fb
S: +OK maildrop has one message (369 octets)

In the example, the MD5 algorithm is applied to the string

<1896.697170952@ima.com>tanstaaf

where “tanstaaf” is the shared secret. By applying the MD5 algorithm to the string, it produces a digest value of

c4c9334bac560ecc979e58001b3e22fb

It must be noted, however, that it is not possible to use the APOP and PASS commands simultaneously. For a given mailbox, either the USER/PASS command sequence or the APOP command must be used, but not both.

SCALING AND OPERATIONAL ISSUES

It has been found that when some of the optional commands are used in large-scale commercial post office operations where users are unrelated, the combination of using

Copyright © 1998 IMA Philippines, Inc. All rights reserved

the UIDL command and not issuing the DELE command can result in a weak version of the “maildrop as semi-permanent repository” functionality normally associated with the Internet Mail Access Protocol (IMAP). Under this setup, already-read messages tend to accumulate on the server without bound. The inability of POP3 to handle maildrops with hundreds or thousands of messages aggravates the situation.

Thus, it is recommended that operators of large-scale multi-user servers impose a per-user maildrop storage quota and/or enforce a policy regarding mail retention on the server. Servers using the first option, however, must make sure that users are informed of the exhaustion of the quota. Server operators using the second option must see to it that users are aware of all the message deletion policies in force.

CONCLUSION

POP3 is an important component of Internet electronic mail. While SMTP is designed to put mail into the mailbox from the Internet, POP3 is designed to retrieve mail from the mailbox and deliver them to their recipients.

However, POP3 also have its shortcomings. With POP3, users need to download their mail and close server connection before they can read their mail. To get more mail, they have to establish connection with the server again.

In addition, in using POP3, a user’s password is transmitted in the clear over the network. Although an optional POP3 command, APOP, has been developed to solve this problem, very few email clients support this feature.

One alternative technology to POP3 that is steadily gaining popularity in the Internet community is IMAP. Unlike POP3, IMAP can function both as an “offline” and an “online” protocol. IMAP also implements a security feature similar to that provided by the APOP command in POP3.

Copyright © 1998 IMA Philippines, Inc. All rights reserved