

GXS



Information Exchange Interface Programming Guide

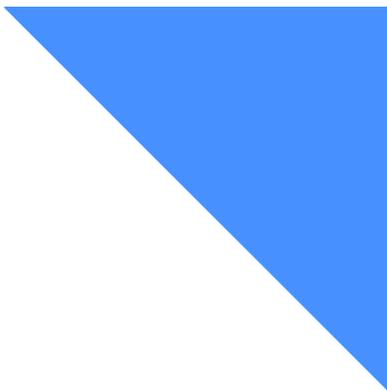
Version 4 Release 6

Eleventh Edition (November 2005)

This edition replaces document number GC34-2222-09.

© Copyright GXS, Inc. 1998, 2005. All rights reserved.

Government Users Restricted Rights - Use, duplication, or disclosure restricted.



Contents

- To the reader ix
 - Who should read this book ix
 - How to use this book x
 - Related books x
 - Expedite/Async xi
 - Expedite Base xi
 - Expedite/CICS xi
 - Expedite/Direct xi
 - Expedite for Windows xi
 - Information Exchange xi

- Chapter 1. Introducing Information Exchange 1
 - What Information Exchange does 1
 - Communication protocols 2
 - Security on the network 3
 - Starting and stopping a session 3
 - Restart and recovery considerations 4
 - Commands, responses, and requests 4
 - Message charges 5
 - Features of Information Exchange 6
 - Information Exchange service parameters 8
 - Archive storage days 8
 - Message limits 8
 - Number of active continuous receive commands 8
 - Output transmission block size 8
 - Storage days for audit, session, receive, and event trails 8
 - Undelivered message retention days 9

- Chapter 2. Controlling the Information Exchange session 11
 - Extended security option 11

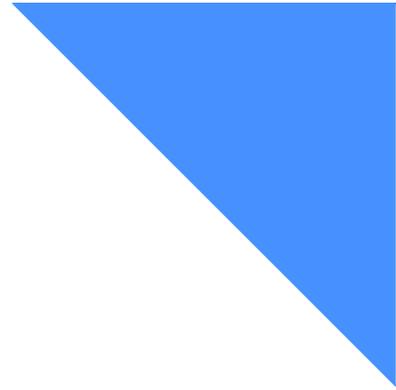
Starting a session	12
How the Session Start command affects receiving messages	13
Session Start response	14
Ending a session	14
Session End response	14
Recovery and restart	14
Message volume and frequency	15
Availability of Information Exchange messages	16
Setting the recovery level	16
Session-level recovery	16
Stream-level recovery	17
Checkpoint-level recovery	18
Transmitting messages	18
Restarting transmission	19
Receiving messages	19
Restarting a receiving operation	20
Group-level recovery	20
Message-level recovery	20
SNA protocols and higher levels of recovery	20
Chapter 3. Information Exchange message services	23
Message definition	23
Message headers	23
Common data header (CDH)	23
CDH format	24
Creating the CDH	24
Receiving the CDH	24
CDH errors	25
Transmitting large messages using message segments	25
Creating message groups	27
Defining a message group	28
Processing carbon copy requests	29
Non-candidates for carbon copy processing	31
Scenarios for carbon copy processing	32
Scenario 1	32
Scenario 2	34
Scenario 3	36
Canceling messages	38
Browsing information about inbound messages	38
Requesting test messages	39
Requesting acknowledgments	39
Obtaining status information	40
Tracking messages and sessions	40
Message audit feature	41
Session trace feature	42
Asynchronous background processing	42
Chapter 4. Sending messages	43
Using the Send Message command	43
Using the force selective receive search criteria	44
Message delivery classes	44
Addressing messages	45

Information Exchange addresses	45
Intersystem Information Exchange addresses	45
Alias addressing	47
Distribution lists	47
Verifying a distribution list	48
Validating addresses, payment levels, and authorizations	48
Organizing data for transmission efficiency	49
Separating data by destination	49
Separating data by type	49
Making user data understandable	50
Self-defining data	50
Data-element splitting	50
Character-code sets	50
Text compression considerations	51
Committing messages	51
Considerations for committing messages	52
Commit response	52
Requesting a pause in transmission—Pace command	52
Pace response	53
Chapter 5. Receiving messages	55
Using the Receive Message command	55
Receiving messages sent using force selective receive search criteria	56
Format of received messages	56
Using the Commit request	56
Responding to a Commit request	57
Understanding the Pace request	57
Responding to a Pace request	57
Messages generated by Information Exchange	57
System addresses	58
Archiving messages	59
Creating groups of archive messages	59
Archiving carbon copy messages	60
Chapter 6. Library services	61
Library functions	62
Defining a library	62
Adding members to the library	62
Deleting members from the library	62
Deleting a library	62
Retrieving a library member	63
Updating library members	63
Browsing library members	63
Listing libraries	64
Listing library members	64
Searching the library	64
Chapter 7. Handling errors	67
Error response message format	68

Chapter 8. Common data header	69
CDH entry fields	69
Chapter 9. Cluster mailbox	77
Cluster user terminology and naming convention	77
Cluster user registration	78
Special considerations for cluster users	78
Handling commands	79
Archiving messages	80
Audit Trails	80
Session/Receive traces	80
Changing the number of cluster mailboxes	80
Chapter 10. Information Exchange commands and responses	81
Commands that cause responses	81
Requests that require responses	82
Summary of command sequences	82
Means of expansion	85
Format of commands and responses	85
Common field name descriptions	85
Finding descriptions	86
Session commands	87
Session Start command—SDISSTA	87
Session Start response	91
Session Start response codes	93
Session End command—SDISEND	96
Session End response	97
Session End response codes	98
Session Inquiry command—SDIINQS	99
Session Inquiry response	100
Long format	100
Short format	101
Session Trace Browse command—SDITRLB	104
Session Trace Browse response	106
Session Trace Retrieve command—SDITRLR	107
Output Format	109
Session Trace Retrieve response	110
Session Trace Receive Record—character format	111
Understanding source address specification	113
Session Trace Record—character format	115
Message transmission commands	118
Cancel command—SDICNCL	118
Understanding an Information Exchange address specification	121
Date processing for expansion level 1	122
Commit command—SDICMIT	123
Commit command response	124
Pace command—SDIPACE	125
Pace command response	126
Probe command—SDIPROB	127
Asynchronous Probe response message text	129
Output format	130
Synchronous Probe response	131

Purge Input command—SDIPURG	133
Send Message command—SDISNDM	134
Long format	134
Short format	135
Understanding destination address specification	141
Mailbox commands	142
Commit request	142
Commit request response—SDICRSP	143
Load Test Messages command—SDILTST	144
Message Inquiry command—SDIINQM	146
Message Inquiry response	147
Message Queue Query command—SDIQUMS	149
Message Queue Query response	151
Message Queue entry	153
Expansion level 1	153
Expansion level 2	153
Expansion level 3	154
Pace request	157
Pace request response—SDIPRSP	158
Purge Message command—SDIPRGM	159
Purge Message response	160
Receive Message command—SDIRCVM	161
Understanding source address specification	168
Received Message header	169
Long format -- Expansion levels 1, 2, and 4	169
Long format -- Expansion level 5	170
Short format	171
Alias commands	176
Alias Inquiry command—SDIINQA	176
Alias Inquiry response	178
Define Alias command—SDIDALS	179
Alias entry	181
Distribution List commands	182
List Define command—SDILSTD	182
G type lists	183
Destination ID entry	185
List Verify command—SDILSTV	186
Output format	188
Library commands	189
Handling library command responses	189
Expansion level 1	189
Expansion level 2	189
Browse Library Member command—SDILBRW	190
Browse Library Member response	192
Define Library command—SDILDEF	194
Define Library response	197
Delete Library command—SDILDEL	199
Delete Library response	200
Delete Library Member command—SDILDLM	201
Delete Library Member response	202
List Libraries command—SDILLST	203
List Libraries response	205

Library Entry item structure	207
List Library Members command—SDILMBR	210
List Library Members response	211
Library Member List entry	213
Retrieve Library Member command—SDILRTV	214
Retrieve Library Member response	217
Retrieve Library Member CDH command—SDIRCDH	219
Retrieve Library Member CDH response	220
Search Library command—SDILSCH	222
Search Library response	224
Library Title command—SDITITL	227
Library Title response	228
Library Member item format	230
Audit Commands	231
Audit Browse command-SDIAUDB	231
Audit Browse response	236
Audit Retrieve command—SDIAUDR	237
Output format	241
Audit Retrieve response	242
Audit Record/Display format	243
Archive commands	258
Archive Retrieve command—SDIARTV	258
Archive Retrieve response	259
Glossary	261
Index	265



To the reader

This book provides the information necessary to create application interfaces for Information Exchange, the mailboxing component of GXS EDI Services. EDI Services, part of the managed electronic transaction services from GXS, is the worldwide market share leader in delivering e-business documents, such as purchase orders and bills of lading, between trading partners. EDI Services provides you with the products, services, and training to transact business electronically, helping you to increase efficiency, eliminate paper waste, and reduce costs.

This book also provides information on how to access and control an Information Exchange session and how to transfer data.

Information Exchange provides a means of sending, storing, and retrieving information electronically, and can make it possible for users on dissimilar computer systems to communicate with one another. By establishing a computer-to-computer communication network between different locations, Information Exchange can both speed and simplify the delivery of messages, electronic data interchange (EDI) envelopes, and other data.

Who should read this book

This book is intended for experienced programmers responsible for creating application interfaces for Information Exchange.

How to use this book

The following typographic conventions are used in this book.

Item	Convention	Example
Command, descriptive name	Initial Caps	Session Start command
Command, actual name	SMALL UPPERCASE LETTERS	SDISSTA
Fields	SMALL UPPERCASE LETTERS	EXPAND
Default values in command formats	Boldface	The default is blank
Field values and variables in command formats	Italics	Valid field values are <i>G</i> , <i>P</i> , and blank
Words in the glossary, first time used in body text	Italics	



NOTE: When blank is listed as a value, it refers to a blank space and not the actual typed word.

Related books

The following books contain information relating to the topics covered in this book. These documents can be viewed on the GXS Web site at <http://www.gxs.com>

Expedite/Async

- *Using Expedite/Async*, GC34-2317

Expedite Base

- *Expedite Base/AIX for RISC System/6000 Programming Guide*, GC34-3280
- *Expedite Base/MVS Programming Guide*, GC34-2204
- *Expedite Base for Windows Programming Guide*, GC34-2253
- *Expedite Base/400 Programming Guide*, GC34-2254

Expedite/CICS

- *Customizing and Developing Applications with Expedite/CICS*, GC34-3304
- *Using the Expedite/CICS Display Application*, GC34-3303

Expedite/Direct

- *Using Expedite/Direct*, GC34-2214

Expedite for Windows

- *Expedite for Windows Software Development Kit Programming Guide*, GC34-3285
- *Expedite for Windows Software Development Kit Programming Reference*, GC34-3284
- *Expedite for Windows User's Guide*, GC34-2341

Information Exchange

- *Information Exchange Administration Mailbox Command Reference*, GC34-2260
- *Information Exchange Administration Services Messages and Codes*, GC34-2323
- *Information Exchange Administration Services User's Guide*, GC34-2221
- *Information Exchange Administration Supplement*, GN10-4073
- *Information Exchange Charges Reference*, GX66-0653
- *Information Exchange Messages and Formats*, GC34-2324
- *Information Exchange via TCP/IP FTP Gateway Quick Reference*, GX66-0664
- *Information Exchange via TCP/IP FTP Gateway User's Guide*, GC34-2345

Introducing Information Exchange

This book provides the information necessary to create application interfaces for Information Exchange, the mailboxing component of GXS EDI Services. EDI Services, part of the managed electronic transaction services from GXS, is the worldwide market share leader in delivering e-business documents, such as purchase orders, and bills of lading, between trading partners. EDI Services provides you with the products, services, and training to transact business electronically, helping you to increase efficiency, eliminate paper waste, and reduce costs.

What Information Exchange does

The Information Exchange mailbox enables communication between dissimilar computer systems. It can link geographically scattered locations of your company or connect you electronically with the trading partners with whom you do business. Contact your network specialist to find out which countries are currently supported.

A unique Information Exchange address containing an account ID and a user ID identifies each Information Exchange user. At the start of each session, this unique address identifies you and your Information Exchange mailbox to Information Exchange.

Information Exchange controls data communication through a variety of commands, making it possible to:

- Send messages
- Cancel messages
- Maintain temporary or permanent lists of addresses
- Use alias addresses
- Receive messages
- Inquire about audit information for messages from and to a particular Information Exchange address
- Inquire about the status of a session

- Inquire about the messages in your mailbox
- Verify the Information Exchange address, payment level, and authorization level of the intended recipients prior to sending messages
- Maintain a library
- Access session trace data

Each command to Information Exchange includes the Information Exchange address of the user on whose behalf the command is to start. The Send Message command identifies where the message will be delivered through the Information Exchange address. When sending to a list of users, the command must give the name of a list of Information Exchange addresses.

With the Information Exchange carbon copy feature, you can send a copy of a message to additional destinations while still sending the message to the original destination. With the redirect feature, you can also redirect a message to an alternate destination.

Your application can communicate with Information Exchange on behalf of one or more Information Exchange addresses. Information Exchange addresses are assigned by the Network during registration, and each Information Exchange address is independent of other Information Exchange addresses.

This means that your program can communicate with Information Exchange for one or more Information Exchange addresses during one physical connection between the systems.

Communication protocols

Information Exchange supports the following communication protocols:

- Synchronous Data Link Control (SDLC), a protocol used with IBM's Systems Network Architecture (SNA) for synchronous information transfer over a link connection
- Binary synchronous (bisynchronous) communication, a protocol for managing synchronous transmission of binary-coded data
- Asynchronous communication (ASC), a protocol used in start-stop communication
- TCP/IP FTP Gateway communication, a protocol that enables Information Exchange to work like an FTP server, allowing you to send and receive EDI transactions through a secure IP network or the Internet

Select the SDLC protocol if you have an SNA system operating with any of the following:

- Information Management System/Virtual Storage (IMS/VS)
- Customer Information Control System/Virtual Storage (CICS/VS)
- Equivalent logical unit (LU) type 6 protocols

The Network supports SDLC through the Advanced Communication Function of the Virtual Telecommunications Access Method (ACF/VTAM*).

SDLC is also appropriate for systems that use LU type 0 or LU type 2 protocols, and systems or terminals supporting LU type 6.2 (Advanced Program-to-Program Communications** or APPC**) protocols.

Bisynchronous systems communicate with Information Exchange through Expedite/Direct or the Batch Data Interface (BDI), which are interfaces resident in the network.

Security on the network

The network Session Access Control application, a program running in the network control processor, directs network and application access through user profiles. Special authorization features, implemented through the user profiles, provide further security.

Network security features operate within a widely used data processing environment. To maximize security protection in this type of environment, you are responsible for securing and changing passwords and profile authorizations at the recommended intervals.

When the network verifies that you are a valid user, you can access Information Exchange if you have registered for this service and if you have entered a valid Information Exchange password. Information Exchange provides its own level of security with both user ID and password security.

The Extended Security Option (ESO) provides additional password and mailbox security. For more information, see “Extended security option” on page 11.

Communication with the TCP/IP FTP Gateway via the Internet is secured by the exchange of authentication certificates and private keys between the FTP client and the Gateway server.

Starting and stopping a session

A session is the time spent communicating with Information Exchange. To start a session, use the Session Start command (SDISSTA). This command identifies you to Information Exchange through your Information Exchange address. To end a session, use the Session End (SDISEND) command. Between session start and session end, use other commands to complete such tasks as transferring messages and retrieving archive messages.

Sessions exist between application programs and Information Exchange. Before sending a Session Start command, you must establish a physical connection between your program and the network by using one of two methods:

- With a leased line, log on to the network.
- Without a leased line, dial the network telephone number and identify yourself to the network.

When you have established communication with the network, send a Session Start command.

Information Exchange can handle sessions for many Information Exchange addresses simultaneously. For this reason, consider programming your application to keep track of data and control various Information Exchange addresses simultaneously.

If your programs cannot handle multiple concurrent sessions, you can use more than one Information Exchange address by managing sessions serially. To do this, complete the current session before starting another session. You do not need to break and re-establish the physical connection between your program and the network.

Restart and recovery considerations

When a communication session is interrupted, restart and recover a session by using one of the following methods:

- Session-level recovery
- Stream-level recovery
- Checkpoint-level recovery
- Group-level recovery
- Message-level recovery

For more information on each recovery level, see “Recovery and restart” on page 14.

Commands, responses, and requests

Application programs must send commands that tell Information Exchange what to do and for whom. When one user sends a message using Information Exchange, the message data is preceded by an Information Exchange message header that serves as a command to Information Exchange.

To receive messages from Information Exchange, issue a Receive Message (SDIRCVM) command. Messages that you receive contain an Information Exchange header, which specifies your

Information Exchange address, and the message text. For more information about the Receive Message command, see “Receive Message command—SDIRCVM” on page 161.

Messages generated as a result of commands you sent to Information Exchange are placed in your Information Exchange mailbox. To receive them, issue a Receive Message command. Such commands include, but are not limited to, the following:

- List Verify
- Cancel
- Load Test Messages
- Probe
- Audit Retrieve
- Session Trace Retrieve

When your application is receiving messages from Information Exchange, Commit or Pace requests are sent by Information Exchange to assist your application in synchronizing the communication. To checkpoint or synchronize messages that you send, use the Commit or Pace command. Requests are identified through the Information Exchange address for which they are issued.

Some commands, such as Session Start, cause Information Exchange to send a response that indicates the results of the command back to your application.

Information Exchange automatically transmits responses as a result of commands. You do not need to issue a Receive Message command to review these responses. These responses are not placed in the user’s mailbox.

Message charges

Everyone who is registered to use Information Exchange has a user profile defining that user's characteristics and ability to perform functions in Information Exchange. You can indicate in your profile whether you are willing to pay for messages sent, received, or both. In part of your profile, the trading partner list, you can list all Information Exchange addresses from which to receive messages and from which not to receive messages. This is accomplished by blocking communication from that address. If you are willing to receive messages from all Information Exchange users, you do not need a list.

Use Information Exchange Administration Services to manage your profile and trading partner list.

Charges for messages that are transmitted via Information Exchange can be allocated to either the message sender or receiver, or they can be split between the two. Message charges are split by dividing the total message path through Information Exchange into the following two domains:

- Send Side

The message path from the sender to Information Exchange. Charges incurred as a result are called send-side charges.

- Receive Side

The message path from Information Exchange to the destination. Charges incurred as a result are called receive-side charges.

In addition, if carbon copy messages are created the following two types of charges may be incurred:

- Carbon copy *processing* charge

The charge associated with processing the carbon copy message.

- Carbon copy *receive* charge

The charge associated with receiving the carbon copy message.

Both charges are incurred for each carbon copy recipient. The processing charge is applied at the time the carbon copy or redirected message is created. This charge is incurred regardless of whether the carbon copy or redirected message is successfully delivered to the recipient's mailbox.

For more information on how message charges work, refer to *Information Exchange Charges Reference*.

Features of Information Exchange

The following is a brief overview of the Information Exchange features you can use. For more information on each feature, see the appropriate chapter.

- **Sending and receiving messages**

The main feature of Information Exchange is to enable you to send and receive messages. Various options specify how you want a message sent or received. For more information, see “Sending messages” on page 43 and “Receiving messages” on page 55.

- **Cluster mailbox**

With the cluster mailbox feature of Information Exchange, you can define a set of mailboxes to appear as one mailbox to your trading partners. This feature allows you to define additional mailboxes for use in sending and receiving data through Information Exchange without requiring your trading partners to change address information. This feature is most beneficial if you send and receive large amounts of data from many trading partners while operating within a restricted window for completing message transport. For more information about the cluster mailbox, see “Cluster mailbox” on page 77.

- **Message cancelation**

With Information Exchange, you can cancel messages that you have already sent, if they have not yet been received at their destination. With the Cancel command, you can cancel one or more messages or message groups, whether you sent them to an individual or to a list of users. For more information, see “Canceling messages” on page 38.

- **Distribution lists**

You can send messages or message groups to several Information Exchange addresses by creating a temporary or permanent distribution list. For more information, see “Distribution lists” on page 47.

- **Message archiving**

You can archive messages for the period of time that you specified in your Information Exchange user profile. You can have all the messages you receive automatically archived (forced archiving), or you can archive only those messages you select by using the archive reference ID (ARCREFID, field 21) of the Receive Message command. For more information, see “Archiving messages” on page 59 and “Receive Message command—SDIRCVM” on page 161.

- **Alias addressing**

With Information Exchange, you can address other users by an alias. Aliases are pseudonyms that you can use instead of an Information Exchange address. These aliases are held in alias tables, which relate the alias to the corresponding Information Exchange address.

When you use aliases to send or receive messages, Information Exchange converts them into Information Exchange addresses for its internal operation. For more information, see “Alias addressing” on page 47.

- **Extended security option**

The Extended Security Option provides additional password and mailbox security. For more information, see “Extended security option” on page 11.

- **Library support**

You can use the library feature of Information Exchange to store data items. Unlike messages stored in your mailbox or archive, the system does not delete information in a library after a period of time has elapsed. Deletion is under the control of the user. With the library feature, you can authorize which users can browse and retrieve data items. For more information, see “Library services” on page 61.

- **Address, payment level, and authorization validation**

Before you send a message, you can validate an Information Exchange address, payment level, and authorization by using the Probe command. For more information, see “Validating addresses, payment levels, and authorizations” on page 48.

- **Message audit**

Information Exchange maintains audit information about the status of your messages. You can browse and retrieve audit information using the audit feature. For more information, see “Message audit feature” on page 41.

- **Session trace**

Information Exchange maintains information about the status of your sessions. You can browse and retrieve session information using the session trace feature. For more information, see “Session trace feature” on page 42.

- **Carbon copy**

The carbon copy feature enables the delivery of an electronic copy (*carbon copy*) of a message to one or more recipients. The carbon copy is sent in addition to the original message. Either the sender or receiver of the message can request a carbon copy to occur. With the carbon copy feature, the receiver can also *redirect* a message to another mailbox. Redirect differs from carbon copy in that the original message is not delivered to the original target mailbox. Instead, the message is purged at the original target mailbox and diverted (redirected) to an alternate mailbox. For more information about carbon copy messages, see “Processing carbon copy requests” on page 29.

Information Exchange service parameters

Information Exchange uses parameters to allow customizing the service across service providers worldwide. The value of these parameters is established by the particular Information Exchange service provider and may vary at the different installations throughout the world.

The following is an alphabetical list of these parameters.

Archive storage days

- The maximum number of days a message can be archived.
- The default number of archive days to be set in the user profile.

Message limits

To protect user applications from looping and to prevent user sessions from timing out during a message commit or restart backout process, Information Exchange imposes two limits upon the user application. The limits are:

- A limit on the size of a message group (sometimes referred to as a message or a file).
- A limit on the number of uncommitted message groups (sometimes referred to as messages or files). Usually this only concerns users that employ session or stream recovery levels.
 - When sending messages to Information Exchange, the user session abnormally terminates when the limit is exceeded.
 - When receiving messages from Information Exchange, the system delivers all messages up to the established limit; then an end-of-mailbox condition is simulated. If the mailbox had messages in excess of the limit, only messages up to the limit are delivered and the rest remain in the mailbox until the next receive command is issued following the committal of the received messages.

Number of active continuous receive commands

A service provider may limit the number of active continuous receive commands a user can have at any given time. Some service providers may choose not to impose a limit.

Output transmission block size

These parameters relate to the transmission block size to be used when messages are transmitted from Information Exchange to the user.

- The maximum transmission size allowed.
- The default transmission size to be used if not specified by the user application.

Storage days for audit, session, receive, and event trails

- The minimum number of days the audit trails and session/receive traces are stored.
- The maximum number of days the audit trails and session/receive traces are stored.

Undelivered message retention days

Two service parameters are associated with the number of days to retain a message in a mailbox without receiving or explicitly purging it.

- Maximum retention days that can be specified by the sender when sending a message to Information Exchange.
- The default number of days to be used for retention of messages if not explicitly specified by the user.

See the information on the MSGRETN field in the “Send Message command—SDISNDM” on page 134 for sample usage of these fields.

Controlling the Information Exchange session

This chapter discusses how you can control your Information Exchange session. It includes information on passwords, starting and ending sessions, and restart and recovery procedures.

Extended security option

The Extended Security Option (ESO) provides additional password and mailbox security for users. The service administrator must use Information Exchange Administration Services to specify a user as an Extended Security Option user for Information Exchange. For more information on using Information Exchange Administration Services for ESO, see the *Information Exchange Administration Services User's Guide*.

ESO contains the following security enhancements:

- You must change your ESO password if it is the same as your user ID. If you do not provide a new password on the Session Start command, your Information Exchange session will be unsuccessful.
- The rules for ESO passwords are as follows:
 - Must not contain the user ID as any part
 - Must be at least six characters in length
 - Must contain at least three different characters
 - Must contain a nonnumeric first and last character
 - Must contain at least one non-alphabetic character
 - Must contain at least one alphabetic character
 - Must contain only the valid characters A-Z, 0-9, and special characters # @ and \$
 - Must not contain more than two identical consecutive characters
 - Must be different from the current and five previous passwords
 - Must not contain more than three identical consecutive characters from the previous password

If the new password does not conform to these rules, it is considered invalid, and your Information Exchange session will be unsuccessful.

- An ESO user ID is revoked if you make three consecutive attempts to start an Information Exchange session with an invalid password. Any further attempts to start an Information Exchange session are unsuccessful until your user ID is reset.
- To reset your ESO user ID, your service administrator uses the Information Exchange Administration Services password-reset function. The ESO user must then use the new password on the next attempt to start an Information Exchange session.

Starting a session

Before you can start an Information Exchange session, you must access the network. How you do this depends on the communication protocols you are using. For more information, see “Communication protocols” on page 2.

Use the Session Start (SDISSTA) command to establish an application-level interface with Information Exchange. The Session Start command identifies you to Information Exchange by your Information Exchange address (account ID and user ID) and a corresponding password. It also establishes parameters that govern certain aspects of the session you request. You can request Information Exchange services only after Information Exchange processes the Session Start command and you receive a Session Start response from Information Exchange.

Your application can communicate with Information Exchange on behalf of multiple Information Exchange addresses during a single physical connection. To do this, Information Exchange establishes separate sessions, either concurrently or serially, with your application. Each session is independently managed by both Information Exchange and your system. Access authority, data security, message transfer, and accounting are all maintained individually for each Information Exchange address.

The Session Start command:

- Identifies you to Information Exchange by account ID, user ID, and password
- Enables you to set a new password for the next session
- Establishes session options to determine how Information Exchange manages message recovery
- Passes information between your system and Information Exchange necessary for restarting a failed session

The Session Start command is detailed in “Session Start command—SDISSTA” on page 87. After transmitting this command to Information Exchange, your system must not transmit any further commands or messages for the same Information Exchange address until the Session Start response is received from Information Exchange. The Session Start command is a response-mode command. This means that Information Exchange sends a response message to the Information Exchange address that sent the command.



NOTE: Information Exchange discards any data it receives before it sends the Session Start response.

How the Session Start command affects receiving messages

Before you can submit a Receive Message command, you must start a session with Information Exchange. The Receive Message command uses the values you specify in the following Session Start command fields to process your messages properly:

RRLTYPE (field 9)

This field specifies the recovery level to be applied. If the Session Start command has *C* or *G* in this field (indicating checkpoint-level recovery or group-level recovery), Information Exchange sends Commit requests to your Information Exchange address. This means that Information Exchange sends Commit requests as you receive messages. If the Session Start command has a value other than *C* or *G* in this field (indicating either message-, stream-, or session-level recovery), Information Exchange sends Pace requests to your Information Exchange address. This means that Information Exchange sends Pace requests as you receive messages.

MAXMSGSZ (field 10)

This field specifies the maximum size of a transmission block, including Information Exchange message headers and text sent to the Information Exchange address. However, this field does not include the size of SNA headers or bisynchronous transmission control characters. The system reblocks user messages to fit within the value you specify in MAXMSGSZ before sending message segments.

MAXMSGCP (field 11)

This field determines the maximum number of message segments that can be sent to your Information Exchange address between Commit or Pace requests. Information Exchange can send the Commit request more frequently than this during receiving operations. A Commit request cannot occur between segments of a multiple-segment message. Therefore, with multiple-segment messages, the system can delay a Commit request past the number of segments specified.



NOTE: The number of segments specified between Commit requests or Pace requests is applicable for the session, regardless of the number of commands you issue for your Information Exchange address.

COMMCMD (field 12)

This field is the command code that Information Exchange must place in the first 8 characters of Commit requests or Pace requests it sends to an Information Exchange address. If RRLTYPE is *C*, Information Exchange sends a Commit request; otherwise, it sends a Pace request.

The Receive Message command determines what is received and establishes the first 8 characters of each message. For more information, see “Receive Message command—SDIRCVM” on page 161, “Commit request” on page 142, and “Pace request” on page 157.

Session Start response

Information Exchange sends a Session Start response in response to the Session Start command. This acknowledges the start of a session or, in case of a failure, tells you why the session did not start. See “Session Start response” on page 91 for the format of the Session Start response and “Session Start response codes” on page 93 for possible Session Start response codes values.



NOTE: If you are using checkpoint-level recovery or group-level recovery and Information Exchange returns the Session Start response with a *G* as the first character of the session access key (SESSKEY - field 4), Information Exchange sends either message text or a Commit response following this Session Start response.

Ending a session

To end a session, send the Session End command to Information Exchange. Any other procedure you use to end a session is an abend (abnormal end). If you detect a problem during transmission, you can end a session abnormally to prevent your message from being committed. To request an abnormal session end, place a value other than zeros or blanks in SENDCODE (field 6) of the Session End command. To restart the session, you must use the Session Start command. Information Exchange also generates an abend if any of the following conditions exist when it processes the Session End command:

- The Receive Message command is incomplete.
- Checkpoint-level recovery is active, and one or more messages have been received by Information Exchange since the last Commit command was received.
- An input multiple-segment message is incomplete.



NOTE: A message group is ended by:

- The last message in the group (*L* in MSGGRPIN)
- The start of another message group (*S* in MSGGRPIN)
- A single message (*blank* in MSGGRPIN)
- An ENDDATA message (*E* in MSGGRPIN)

See “Session End command—SDISEND” on page 96 for the format of the Session End command.

Session End response

Information Exchange sends a Session End response in response to the Session End command. This response indicates the results of processing the Session End command. For more information, see “Session End response” and “Session End response codes” on page 98.

Recovery and restart

Information Exchange provides five methods of recovery and restart of a logical session. Four of these permit restart and recovery between the end-to-end application code in Information Exchange and your system. A fifth method, message-level recovery, is also provided. This method does not use an end-to-end application protocol, but relies on the recoverability of SNA sessions.

All five methods use message committal. A message is committed when it is transferred from the sender to Information Exchange, and Information Exchange accepts it. A message is also committed when it is transferred from Information Exchange to the receiver, and the receiver accepts it. The primary difference between the five methods is the time when a message committal takes place:

- Session-level recovery

Messages are committed at the normal end of an Information Exchange session.

- Stream-level recovery

Input messages are committed when Information Exchange receives an ENDDATA message. Output messages are committed when a session ends normally.

- Group-level recovery

Messages are committed periodically during a session as a result of specific exchanges. Messages are only committed between message groups during receive message processing.

- Checkpoint-level recovery

Messages are committed periodically during a session as a result of specific exchanges.

- Message-level recovery

When sending messages to Information Exchange, the message is routed to its destination when the last piece is received by Information Exchange. When receiving messages from Information Exchange, the message is deleted from the mailbox when the last piece is sent to the user.

Message volume and frequency

Consider the volume and frequency of message traffic you expect to handle with Information Exchange. If you expect low-traffic volume and only one or two relatively short sessions each day with Information Exchange, use session- or stream-level recovery. In this situation, restarting the session from the beginning costs little extra time and line or computer resources.

If you expect your application to be in almost continuous communication with Information Exchange or if you handle very large volumes of messages, use checkpoint-level recovery.

Availability of Information Exchange messages

You cannot retrieve error and warning messages that Information Exchange detects during Send Message processing until the system commits the associated input messages. Therefore, the availability of these error or warning messages depends on the recovery level you are using. They are first available for retrieval as follows:

For this recovery level:	Error or warning messages are available:
Session-level	In a subsequent Information Exchange session.
Stream-level	Following the submission of an ENDDATA message to Information Exchange.
Group-level or Checkpoint-level	After the Commit command and response are exchanged following the completion of the current message group.
Message-level	As soon as the individual message or message group is received by Information Exchange.

Setting the recovery level

To set the recovery level for a session, use a value in RRLTYPE (field 9) of the Session Start command, as follows:

For this level:	Use:
Session-level	X
Stream-level	S
Group-level	G
Checkpoint-level	C
Message-level	M

If the previous communication session with Information Exchange fails, you cannot change the recovery level during a subsequent session restart, unless you specify one of the reset session options in TESTREST (field 13) of the Session Start command. However, after you have received a normal Session End response, you may specify a different recovery level in the next Session Start command (because this is a new session, not a restart).

A normal session end occurs when the value of SENDCODE (field 5) in the Session End response is zero.

For more information, see “Session Start command—SDISSTA” on page 87 and “Session End command—SDISEND” on page 96.

Session-level recovery

Information Exchange commits all input or output messages when a session ends normally; that is, when the value in SENDCODE (field 5) of the Session End response is zero.

You may leave the SNDCKPTN and RCVCKPTN (fields 14 and 15) of the Session Start command blank; you do not need to maintain checkpoint numbers for session-level restart.

Examine the reply to the Session Start request to determine the status of the previous session and whether to perform the appropriate restart procedures.

The value in RESPRSN (field 5) of the Session Start response indicates the results of processing the Session Start command. See “Session Start response codes” on page 93. If RESPRSN indicates that the Information Exchange session is restarting, Information Exchange purges all messages sent to it during the failed session. However, you can receive messages sent by Information Exchange by submitting the appropriate Receive Message command. The results of restarting with session-level recovery are as if the previous session never occurred. You must resubmit any temporary distribution lists that have been defined and any Receive Message commands that were active, because the system eliminates them.

With session-level message recovery, LSTSNDP and LSTRVCP (fields 6 and 7) of the Session Start response always contain zeros, indicating that no check-points are in use. You must retain the ability to resubmit all data messages until the receiver has received and processed the data. With session-level recovery, however, you need not keep track of any checkpoint identification.

Restart the session as follows:

1. Discard all messages received during the previous session.
2. Redefine all temporary distribution lists you are using.
3. Resubmit all input messages.
4. Resubmit all Receive Message commands.

Stream-level recovery

Using stream-level recovery, you can handle a single stream of input and a single stream of session output independent of each other, so they can be recovered separately. This method uses separate checkpoints for input and output, and requires the input stream to be complete before the output stream is requested.

If you use stream-level recovery, you must submit all input data messages as a single stream of requests, delimited by an ENDDATA message. This message signifies the end of data input and tells Information Exchange to commit the previous input messages and to queue them to their destinations.

If the session ends while the input stream is incomplete, Information Exchange discards all messages in the input stream. If you do not resubmit the messages after restarting the session, they are lost.

Information Exchange indicates that processing input messages is complete by:

- A value of zero in SENDCODE (field 5) of the Session End response.
- A value of zero in LSTSNDP (field 6) in a subsequent Session Start response if communication fails during the processing of the Session End command.

Information Exchange indicates that it discarded an input message stream by placing a value of 1 in LSTRVCP of the subsequent Session Start response.

Messages you receive during the session are committed (deleted from Information Exchange) when you submit a normal Session End command following the completion of one or more output requests. An output request is complete when an ENDDATA message is received for the corresponding Receive Message command. Information Exchange then indicates it has committed (deleted) the messages you have received by:

- Transmitting a value of zero in SENDCODE in the Session End response for the current session.
- Transmitting a value of zero in LSTSNDP in the following Session Start response if communication fails during the processing of the Session End command.

Examine the Session Start response to determine whether you should perform the recovery procedures. If RESPRSN (field 5) of the Session Start response is less than 64, indicating the session started, you must also examine LSTSNDP and LSTRCVCP (fields 6 and 7) to determine whether you need to perform the restart procedures, as follows:

- If LSTSNDP contains a value other than zero, Information Exchange discards all input data messages submitted during the most recent input stream of the previous session. See “Session Start response codes” on page 93, and resubmit the messages.
- If LSTRCVCP contains a value other than zero, all output that was received from Information Exchange during the previous session remains to be retransmitted. Discard output received during the previous session and resubmit the Receive Message commands.



NOTE: The LSTSNDP and LSTRCVCP values reflect the status of the most recent session.

When restarting stream-level sessions, Information Exchange retains any temporary distribution lists defined during the previous aborted session.

You must retain the ability to resubmit all input data messages of a session until Information Exchange indicates they have been committed, either by:

- A value of zero in SENDCODE in the current Session End response.
- A value of zero in LSTSNDP in the next Session Start response.

With stream-level recovery, you need not keep track of any checkpoint identification.

Checkpoint-level recovery

Information Exchange implements checkpoint-level recovery as a two-stage commit process using Commit commands (or requests) and responses. When Information Exchange processes the Commit command, it assigns a checkpoint number for message-input operations. When you process the Commit request, you must assign a checkpoint number for message-output operations. These numbers are used in the Session Start command and response to start recovery at a predetermined point in the session. In this method, Information Exchange handles the flow of messages in each direction separately. The input and output checkpoints are independent of each other.

After sending one or more messages to Information Exchange, send the Commit command to request that the messages be committed. Do not send a Commit command between segments of a multiple-segment message. When Information Exchange transmits messages to you, it asks you to commit messages by sending you a Commit request.

Transmitting messages

To transmit messages with checkpoint-level recovery:

1. Transmit one or more messages.

2. Transmit a Commit command and wait for the Commit response.
3. Check the checkpoint response code in `CHKPNTCD` (field 6) of the Commit response. A value of 8 or greater indicates that a severe error occurred on the previous set of send messages. This is likely to occur only during your initial testing. Check the headers on the messages sent to Information Exchange since the last commit; also, check the fields on the Commit command itself. To receive any Information Exchange error messages that the system generates, proceed to the next commit point. You can do this by sending another Commit command, or by sending a message and issuing a Receive Message command.
4. If `CHKPNTCD` is 4 or less, delete the transmitted messages from your transmit storage so they are not retransmitted when you restart the session.
5. Retain the checkpoint number in `CHKPNTNO` (field 5) of the Commit response in case you need it to restart the session.
6. Resume message transmission.

Restarting transmission

To restart message transmission:

1. Place the retained checkpoint number in `SNDCKPTN` (field 14) of the Session Start command.
2. Check the value of `RESPRSN` (field 5) of the Session Start response and be sure that it is less than 64. If it is 64 or greater, the session has not been restarted.
3. If `RESPRSN` is less than 64, check `LSTSNDCP` (field 6) of the Session Start response to see if it is the same as the retained checkpoint number you placed in `SNDCKPTN`; they should always be the same.
4. Resume message transmission with the first available message in your transmit storage.

Receiving messages

To properly process messages with checkpoint-level recovery, you must have storage space to hold the messages you receive until they are committed. When they are committed, place them into the storage area where they are processed. There is a checkpoint number associated with the receive storage that you must increase on each receive-commit exchange. Keep this number in a safe storage place, such as on a disk.

The basic procedure to receive messages is as follows:

1. As you receive each message, place it in the receive storage.
2. When you receive the Commit request, increase the checkpoint number and place it in `CHKPNTNO` (field 5) of the Commit response.
3. Send the Commit response.
4. Move the messages from the receive storage into the processing storage.
5. When you receive the next message, place it in the receive storage.

Restarting a receiving operation

With Information Exchange, you can recover if a session fails when the commit exchange is in progress.

To restart the receiving operation:

1. Place the checkpoint number in RCVCKPTN (field 15) of the Session Start command.
2. Ensure the value in LSTRCVCP (field 7) of the Session Start response is the same as the value in RCVCKPTN.
3. Delete all messages in the receive storage because Information Exchange retransmits them.



NOTE: If messages were being sent to your system when the previous session failed, Information Exchange begins to transmit messages immediately following the Session Start response. With checkpoint-level recovery, the session is restored to its previous state. You do not need to resend either the Receive commands or the temporary distribution lists.

Group-level recovery

Information Exchange implements group-level recovery in the same manner as checkpoint-level recovery. The only difference is that Information Exchange will not send commit requests between message groups during message receive processing.

Message-level recovery

Implementation of message-level recovery depends on the message-recovery abilities of SNA because Information Exchange uses the SNA set-and-test-sequence-number (STSN) protocol to synchronize the flow of messages when it restarts a physical session.

With this recovery level, when Information Exchange receives the last segment of a message (or the last segment of the last message in a message group), the system routes the message (or message group) to its indicated destination.

Similarly, when Information Exchange sends you the last segment of a message (or the last segment of the last message of a message group), it deletes the message (or message group) from Information Exchange.

SNA protocols and higher levels of recovery

If you are an SNA user who has implemented session-, stream-, group-, or checkpoint-levels of recovery, you must take a special precaution. Because the SNA products provide a message resynchronization process and Information Exchange provides a session resynchronization process, you must follow a certain procedure to restart a session.

Normally, Information Exchange sends output to you immediately after the SNA physical session is re-established. If you are using one of the higher levels of recovery, such as session-, stream-, group-, or checkpoint-level, ignore all output from Information Exchange and send the Session Start command.

The Session Start response should be the first output from Information Exchange that the system processes after you restart the SNA session. You can ensure that this happens by using a switch in your program. If the switch is off (its initial condition), all output you receive from Information Exchange is inspected to see if it is your Session Start response. If the output is not the expected

Session Start response, the message you have just received is ignored (discarded). If the output you received is the Session Start response, set the switch on and the system processes your Session Start response. The system processes subsequent output from Information Exchange according to its type.

Because of the Session Start command, Information Exchange ignores any intervening input it receives (except Commit commands) between the SNA session restart and the arrival of the Session Start command.

Information Exchange message services

This chapter discusses Information Exchange message services. Message services assist you in transferring and processing your messages. This chapter also includes information on how to use the common data header (CDH) and the message and session tracking features.

Message definition

A message is any piece of data that a computer can transmit electronically to another computer. Messages are transmitted to or from Information Exchange in an envelope. Each message envelope contains a header and the actual message data. When you send a message to Information Exchange, the Send Message command contains both header and address information.

Message headers

A message header is the beginning of a message, prior to the text of the actual message. A message header contains command and address information. When you send a message to Information Exchange, the header indicates your address and the address of the intended recipient. When you receive a message, the header provides information about the sender.

Common data header (CDH)

To sufficiently handle the data it receives, a receiving system often requires more information about a message than a normal message header can supply. Your application can use a CDH to send and receive additional information that describes the contents of a message. A typical CDH includes:

- The type of message (text or binary)
- The record delimiters used in the message (for example, carriage-return and line-feed (CRLF) characters or record lengths)
- Information on whether the data is formatted for electronic data interchange (EDI) and the type of EDI data (for example, X12)

- A free-format text description
- The original name of the file on the sending system

The CDH must be the first message of a message group. To indicate the CDH, use MSGNCLS (field 13) of the Send Message command. Specify an **S** in this field if you are sending a normal-mode message with a CDH. Specify an **R** in this field if you are sending a test-mode message with a CDH. For more information, see “Creating message groups” on page 27 and the “Send Message command—SDISNDM” on page 134.

CDH format

The CDH always begins with a 2-byte field (HLENGTH) that describes the length of the entire header, including the length field itself. The CDH is comprised of variable-length entries that are order-independent. Each entry contains a 2-byte length field and a 1-byte ID, followed by the variable-length data. The entire CDH cannot exceed 32K bytes in length. The entries of the CDH are described in “Common data header” on page 69. The following illustration shows an Information Exchange message group in CDH format.

Total Header Length (2 Bytes)	Entry 1 Length (2 Bytes)	Entry 1 ID (1 Byte)	Entry 1
Entry 2 Length (2 Bytes)	Entry 2 ID (1 Byte)	Entry 2	...
Entry N Length (2 Bytes)	Entry N ID (1 Byte)	Entry N	

Figure 1. CDH format. First message of an Information Exchange message group

Creating the CDH

Create a CDH for every message and file that you send from your application. Set those fields of the CDH that apply to your application. Information Exchange assumes a default value for some fields that are not set. See “CDH entry fields” on page 69 for the CDH default values.

Receiving the CDH

If present, you can request Information Exchange to send the CDH, by specifying a 2 in EXPAND (field 5) of the Receive Message command. If you specify a 1 in EXPAND, Information Exchange does not send a CDH to your application even if one is present. Information Exchange sends a CDH only if the sending application created one.

Your application should support every field in an incoming CDH. However, some fields might not be present in an incoming header, because not all sending applications can support each field or the sending application elected to take the default for a field. If a supported field is not present in a received header, assume the default value.

Future Information Exchange releases can continue to add fields to the CDH to support new features. For this reason, it is necessary for applications to handle new and unknown header entries. Program your application to ignore unknown header entries.

CDH errors

There are times when your application might receive headers that have missing or unknown entries. Because these situations are often due to the differences in the release levels of Expedite interface packages, Information Exchange does not consider them to be errors. However, if your application receives an erroneous header, it should be able to interpret the problem and notify the receiving user. The following illustrates possible error conditions and appropriate application responses:

For this error condition:	Program your application to:
The length of an entry is longer than the maximum value.	Truncate the entry at the maximum value.
The value of the entry is invalid.	Use a default value for the entry.
There are duplicate entries.	Take the value of the last entry. If the entry is invalid, use a default value.
The length of the CDH exceeds the length of the Information Exchange message.	Process only to the length of the message.

Transmitting large messages using message segments

The amount of data you can transmit during an Information Exchange session depends on your computer system. Information Exchange permits you to transmit a maximum of 26,000 characters in one physical block. You must include all data and the transmission control characters within this block. However, Information Exchange accommodates messages that exceed this transmission capacity by using message headers to divide long messages and transmit them as smaller segments. Information Exchange uses two types of message header formats: long and short.

A *long header* contains all the necessary routing and control information. Use the long-format header when sending the first segment of a multiple-segment message.

A *short header* contains enough data to relate the message text to the information in the long header. Use the short-format header for subsequent segments.

MSGTXTOV (field 5) of the Send Message command indicates which header and message segment are being used. Valid values for MSGTXTOV (field 5) are the following:

This value:	Indicates:
Blank	The only segment of a single-segment message. Use the long-format header.
S	The first segment of a multiple-segment message. Use the long-format header.
C	An intermediate segment of a multiple-segment message. Use the short-format header.
E	The last segment of a multiple-segment message. Use the short-format header.

For more information, see “Send Message command—SDISNDM” on page 134.



NOTE: If you are using SNA communication, you cannot create message segments larger than the amount of data your system can send in a single transmission (RU chain). If you are using bisynchronous communication, you must not create segments larger than a single bisynchronous transmission record unless you use record chaining.

If you issue a Receive Message command and the message you are about to receive is larger than your system can accommodate, Information Exchange segments the message automatically according to the value you specify in MAXMSGSZ (field 10) of the Session Start command. When starting a session, you must specify the maximum amount of data your system can handle in one physical block.

Message segments and their headers enable different users to send and receive different sizes of transmission units. When Information Exchange receives messages from an Information Exchange address in multiple segments, it concatenates the data in these segments and removes the short-format headers as the system stores the data. Information Exchange then sends the message to its destination by splitting the data into message segments that fit the destination's transmission unit size and inserting the short-format headers.

Users with different transmission unit sizes will see the message segmented differently. However, the continuous string of message-text characters are identical regardless of how Information Exchange segments the message.

The following rules apply to multiple-segment messages:

- Multiple-segment messages can occur within a message group.
- All segments of a multiple-segment message must be sent without interruption.



NOTE: Do not send a Pace or Commit command between segments of a multiple-segment message.

- Use a long-format header for the first or only segment of a message.
- Use a short-format header for the second and subsequent segments of a message.
- The final segment of a multiple-segment message must contain a short-format header with MSGXTOV equal to *E*.

A segmented message is illustrated in Figure 2.

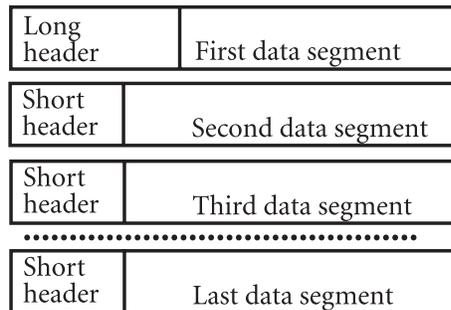


Figure 2. Multiple-segment message

A message can be transmitted in as many segments as necessary. Consider the first character of text following a short header to be concatenated to the last character of text in the previous segment.

Creating message groups

You can combine messages into larger entities called message groups. Information Exchange processes message groups as single complete entities, so other messages do not interrupt the transmission. Message groups enable you to transmit very long transactions and still maintain timely recovery using checkpoints or commits between messages within the group.

For example, a group may consist of several messages, each comprising one page of a report. The report is not complete until Information Exchange transmits all messages, because any single message cannot convey the full meaning of the report. By using the Information Exchange message-grouping facility, you can ensure that individual messages are delivered in the correct order, without interruption by message traffic from other users.

A message group can contain any number of messages of any size. Multiple-segment messages can exist within a message group. The diagram on the next page illustrates the structure of a message group.

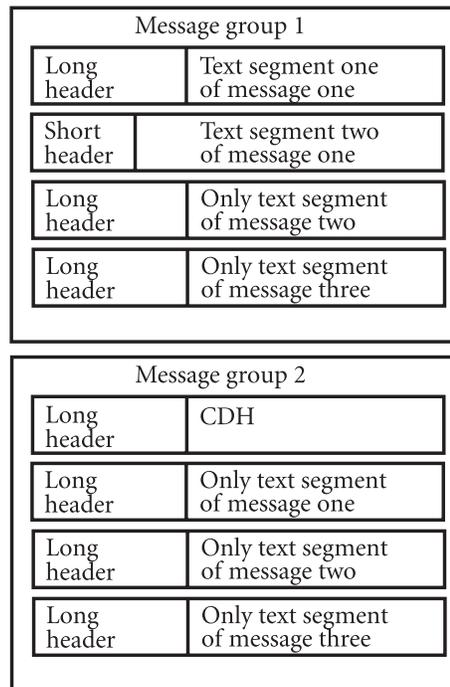


Figure 3. Message group

In Figure 3, two message groups are shown. The first has three records, and the second has three records and a CDH. The first record of the first group is contained in two segments; the other records are only single-segment records.

Defining a message group

To define a message group, use the grouping indicator in MSGGRPIN (field 7) of the Send Message command. To indicate the start of a group, specify an **S** in MSGGRPIN. To indicate subsequent messages within the group, specify an **C** in this field. Specify an **L** to indicate that the current message is the last message in the message group. When you end a message group in this manner, it is called explicit group termination. A group is implicitly terminated by a message header with an *S*, *blank*, or *E* in this field. You can end a message group using explicit or implicit termination only. You must end a group before you submit a Session End command to end a session. A request for abnormal session end with a Session End command ends the session abnormally, but does not end a message group if you are using checkpoint- or message-level recovery. In this case, you must resume transmission of the message group when you restart the session.

When Information Exchange processes a message group, it inherits the following attributes from the first message in the group:

- The destination of the first message in a group determines the destination of the entire group. Subsequent messages within the group must specify the same value as the first message in the following fields of the Send Message command. Otherwise, Information Exchange rejects the entire group.
 - DESTACCT (field 8)
 - DESTUID (field 9)
 - DESTTYPE (field 10)
 - DTBLTYP (field 22)
 - DTBLID (field 23)
- The name field of the first message establishes a name for the entire group. However, individual messages within the group can contain their unique message names.
- The first message of the group also establishes these values:
 - MSGUCLS (field 14)
 - MSGCLASS (field 15)
 - MSGRCPTS (field 16)
 - MSGCHRG (field 17)

See the Send Message command in “Send Message command—SDISNDM” on page 134 for a description of the fields.

Processing carbon copy requests

The carbon copy function enables the delivery of an electronic copy (*carbon copy*) of a message to one or more recipients. The carbon copy is sent in addition to the original message. Either the sender or receiver of the message can request a carbon copy to occur. With the carbon copy function, the receiver can also *redirect* a message to another mailbox. Redirect differs from carbon copy in that the original message is not delivered to the original target mailbox. Instead, the message is purged at the original target mailbox and diverted (redirected) to an alternate mailbox.



NOTE: The original message must be deliverable in order for carbon copy or redirect processing to occur.

- Processing carbon copies requested by the receiver

A message is a candidate for carbon copy processing during the routing of the message to the mailbox of the intended receiver (requester) if the receiver of the message has receive-side carbon copy relationships defined for the message sender (partner).

If the intended receiver has receive-side relationships defined and none specify that the original message be purged:

- The original message is delivered to the original intended receiver.
- Information Exchange determines if any carbon copies should be created.
- A list of carbon copy recipients is generated and each message is created and routed to the carbon copy recipient’s mailbox.

If the intended receiver has receive-side relationships defined, and at least one relationship specifies that the original message be purged:

- The original message is not immediately delivered to the intended receiver.
- Information Exchange determines if any carbon copies should be created.
- A list of carbon copy recipients is generated and each message is created and routed to the carbon copy recipient's mailbox.

If a carbon copy was created as a result of a relationship that specified to purge the original message (for only the first carbon copy generated):

- The original message is routed (redirected) to the carbon copy recipient's mailbox instead of the original intended receiver.
- A delivery acknowledgment is generated for the original message if the original sender requested one.

If no carbon copy was created as a result of a relationship that specified to purge the original message:

- The original message is delivered to the original intended receiver.

■ Processing carbon copies requested by the sender

A message is a candidate for carbon copy processing during the routing of the message to the mailbox of the intended receiver (partner) if the sender (requester) of the message has send-side carbon copy relationships defined.

- Information Exchange determines if any carbon copies should be created.
- A list of carbon copy recipients is generated and each message is created and routed to the carbon copy recipient's mailbox.

■ Routing of carbon copy or redirected messages

A carbon copy or redirected message will be delivered if both of the following are true:

- The recipient is a valid user.
- The carbon copy receive charge is resolved



NOTE: The processing charge is applied at the time the carbon copy or redirect message is created. This charge is incurred regardless of whether the carbon copy or redirect message is delivered to the recipient.

Refer to the *Information Exchange Administration Services User's Guide* for more details on the creation of carbon copy relationships, carbon copy processing, and restrictions.

Non-candidates for carbon copy processing

Carbon copy requests will not be processed for the following:

- When messages are retrieved from archive.
- When library members are retrieved into a user's mailbox.
- When the account ID of the partner is *SYSTEM*.
- When the message is already a carbon copy /redirect message.

Scenarios for carbon copy processing

The following scenarios are provided to illustrate carbon copy processing.

Scenario 1

User A (SEND1.USERA) sends a non-EDI message to User 1 (RECV1.USER1) and requests delivery acknowledgment.

The following relationships have been defined for the sender and intended receiver. The relationships specify all message types and have no criteria specified.

ID	CC Requester	CC Type	Partner	Recipient	Purge Original?
1	RECV1.USER1	Receive	SEND1.USERA	RECIP1.USERX	Y
2	RECV1.USER1	Receive	SEND1.USERA	RECIP1.USERY	N
3	RECV1.USER1	Receive	SEND1.USERA	RECIP1.USERZ	Y
4	RECV1.USER1	Receive	SEND2.USERB	RECIP1.USERW	N
5	SEND1.USERA	Send	RECV1.USER1	RECIP1.USERF	n/a
6	SEND1.USERA	Send	RECV1.USER1	RECIP1.USERG	n/a
7	SEND2.USERB	Send	RECV1.USER1	RECIP1.USERH	n/a

The following outlines the carbon copy processing steps for scenario 1.

1. Check receive-side relationships for intended receiver.
 - a. Determine if **purge original = y** is specified on any receive-side relationship for the intended receiver.

Does the intended receiver have receive-side relationships defined and at least one relationship specifies **purge original = y**?

Answer: Yes, relationship IDs 1 and 3 specify **purge original = y**.

Action: Original message is not immediately delivered.
 - b. Match receive-side relationships to message and generate recipient list.

Determine which carbon copy relationships defined for the receiver match the message and generate a list of carbon copy recipients.

Action: Generate a list of carbon copy recipients, noting if original message should be purged:

Recipient 1: RECIP1.USERX, purge = y

Recipient 2: RECIP1.USERY, purge = n

Recipient 3: RECIP1.USERZ, purge = y

- c. Determine if **purge original = y** on the list of recipients.

Do any of the carbon copy recipients have **purge original = y** specified?

Answer: Yes, Recipients 1 and 3.

Action: Mark original message as redirected to the first recipient with **purge original = y** specified, Recipient 1 (RECIPI.USERX). Send a delivery acknowledgment to SEND1.USERA for original message.

- d. Create and route carbon copy messages.

Create and route carbon copy messages to Recipients 1, 2, and 3.

Action: Carbon copy messages generated and routed to the following recipients:

RECIPI.USERX, purge = y (Message type: Redirected)

RECIPI.USERY, purge = n (Message type: Receive carbon copy)

RECIPI.USERZ, purge = y (Message type: Receive carbon copy)

2. Check send-side relationships for sender.

- a. Match send-side relationships to message and generate recipient list.

Determine which carbon copy relationships defined for the sender match the message and generate a list of carbon copy recipients.

Action: Generate a list of carbon copy recipients.

Recipient 1: RECIPI.USERF

Recipient 2: RECIPI.USERG

- b. Create and route carbon copy messages.

Create and route carbon copy messages to Recipients 1 and 2.

Action: Carbon copy messages generated and routed to the following recipients:

RECIPI.USERF (Message type: Send carbon copy)

RECIPI.USERG (Message type: Send carbon copy)

Scenario 2

User B (SEND2.USERB) sends a non-EDI message to User 2 (RECV2.USER2) and requests delivery acknowledgment.

The following relationships have been defined for the sender and intended receiver. The relationships specify all message types and have no criteria specified.

ID	CC Requester	CC Type	Partner	Recipient	Purge Original?
1	RECV2.USER2	Receive	SEND2.USERB	RECIP1.USERW	N
2	RECV2.USER2	Receive	SEND1.USERA	RECIP1.USERX	Y
3	SEND2.USERB	Send	RECV2.USER2	RECIP1.USERF	n/a
4	SEND2.USERB	Send	RECV2.USER2	RECIP1.USERG	n/a
5	SEND2.USERB	Send	RECV1.USER1	RECIP1.USERH	n/a

The following outlines the carbon copy processing steps for scenario 2.

1. Check receive-side relationships for intended receiver.
 - a. Determine if **purge original = y** is specified on any receive-side relationship for the intended receiver.

Does the intended receiver have receive-side relationships defined and at least one relationship specifies **purge original = y**?

Answer: Yes, relationship ID 2 specifies **purge original = y**.

Action: Original message is not immediately delivered.
 - b. Match receive-side relationships to message and generate recipient list.

Determine which carbon copy relationships defined for the receiver match the message and generate a list of carbon copy recipients.

Action: Generate a list of carbon copy recipients, noting if original message should be purged:

Recipient 1: RECIP1.USERW, **purge original = n**
 - c. Determine if **purge original = y** is specified on the list of recipients.

Do any of the carbon copy recipients have **purge original = y** specified?

Answer: No.

Action: Deliver original message to RECV2.USER2.
 - d. Create and route a carbon copy message.

Create and route a carbon copy message to Recipient 1.

Action: Carbon copy message generated and routed to the following recipient:

RECIP1.USERW (Message type: Receive carbon copy)

2. Check send-side relationships for sender.

a. Match send-side relationships to message and generate recipient list.

Determine which carbon copy relationships defined for the sender match the message and generate a list of carbon copy recipients.

Action: Generate a list of carbon copy recipients.

Recipient 1: RECIPI1.USERF

Recipient 2: RECIPI1.USERG

b. Create and route carbon copy messages.

Create and route carbon copy messages to Recipients 1 and 2.

Action: Carbon copy messages generated and routed to the following recipients:

RECIPI1.USERF (Message type: Send carbon copy)

RECIPI1.USERG (Message type: Send carbon copy)

Scenario 3

User B (SEND2.USERB) sends a non-EDI message to User 2 (RECV2.USER2) and requests delivery acknowledgment.

The following relationships have been defined for the sender and intended receiver. The relationships specify all message types and have no criteria specified.

ID	CC Requester	CC Type	Partner	Recipient	Purge Original?
1	RECV2.USER2	Receive	SEND2.USERB	RECIP1.USERW	N
2	SEND2.USERB	Send	RECV2.USER2	RECIP1.USERF	n/a

The following outlines the carbon copy processing steps for scenario 3.

1. Check receive-side relationships for intended receiver.
 - a. Determine if **purge original = y** is specified on any receive-side relationship for the intended receiver.

Does the intended receiver have receive-side relationships defined and at least one relationship specifies **purge original = y**?

Answer: No.

Action: Original message is immediately delivered to RECV2.USER2.

- b. Match receive-side relationships to message and generate recipient list.

Determine which carbon copy relationships defined for the receiver match the message and generate a list of carbon copy recipients.

Action: Generate a list of carbon copy recipients:

Recipient 1: RECIP1.USERW

- c. Create and route a carbon copy message.

Create and route a carbon copy message to Recipient 1.

Action: Carbon copy message generated and routed to the following recipient:

RECIP1.USERW (Message type: Receive carbon copy)

2. Check send-side relationships for sender.

a. Match send-side relationships to message and generate recipient list.

Determine which carbon copy relationships defined for the sender match the message and generate a list of carbon copy recipients.

Action: Generate a list of carbon copy recipients.

Recipient 1: RECIPI1.USERF

b. Create and route carbon copy messages.

Create and route carbon copy messages to Recipient 1.

Action: Carbon copy message generated and routed to the following recipient:

RECIPI1.USERF (Message type: Send carbon copy)

Canceling messages

Use the Cancel command to cancel messages already sent but not yet received by the addressee. The Cancel command consists of optional and required fields. You must complete all required fields for Information Exchange to process the command. Complete only those optional fields that you need to specify the particular messages you want to cancel.

When Information Exchange processes a Cancel command, all messages and message groups that you specified in a message delivery class and sent to a destination are candidates for cancellation. If you do not specify optional fields in the Cancel command, Information Exchange cancels all messages and message groups sent to the destination in the specified message delivery class.

You can limit the scope of the Cancel command by specifying parameters in the optional fields of the command. These fields must match the corresponding fields in the message header of the first (or only) message of any group you want to cancel. Only messages and message groups that are not received by the destination are candidates for cancellation. Once a message or message group has been delivered by Information Exchange, it cannot be canceled. For more information, see “Cancel command—SDICNCL” on page 118.



NOTE: Information Exchange cannot cancel intersystem messages.

Browsing information about inbound messages

You can retrieve and display details of messages and message groups in your Information Exchange mailbox using the Message Queue Query command. This can help you decide which messages you want to receive and how you want to deal with them. To specify a particular message you want to receive, use the message ID (MSGID - field 1) of the Message queue entry. For more information, see “Message Queue Query command—SDIQUMS” on page 149 and “Message Queue entry” on page 153.

Requesting test messages

To ensure that Information Exchange is sending and delivering your messages correctly, use the Load Test Messages command. There are currently six types of test messages available in Information Exchange. Select messages by typing a number between 0 and 5 in LOMSGNO and HIMSGNO (fields 5 and 6) of the Load Test Messages command. Type the number of the first of a series of test messages in LOMSGNO. Type the number of the last of a series of test messages in HIMSGNO. Information Exchange processes your request in the following manner:

- If the number you specify in LOMSGNO is equal to the number in HIMSGNO, Information Exchange delivers only the message number in LOMSGNO.
- If the number you specify in LOMSGNO is less than the number in HIMSGNO, Information Exchange delivers all messages in ascending order.

To request several nonessential messages, submit successive Load Test Messages commands and type the same value in the LOMSGNO and HIMSGNO fields.

To select and receive test messages, you must:

- Submit the Load Test Messages command to select messages.
- Submit the Receive Message command to retrieve queued messages.

To identify test messages, look at DESTACCT and DESTUID (fields 8 and 9) of the Received Message header. DESTACCT contains a value of `*SYSTEM*`, and DESTUID has a value of `*TSTMSG*` for test messages.

To receive test messages, specify `*SYSTEM*` in DESTACCT (field 6) and `*TSTMSG*` in DESTUID (field 7) of the Receive Message command. You can also use `*SYSTEM**TSTMSG*` in a list referenced by the Receive Message command.

For more information, see “Load Test Messages command—SDILTST” on page 144.

Requesting acknowledgments

Although they are not error messages, Information Exchange acknowledgments are designated by an account ID of `*SYSTEM*` and a user ID of `*ERRMSG*`. The following list describes the three types of acknowledgments that you can request by using MSGRCPTS (field 16) of the Send Message command:

This acknowledgment:	Indicates:
RECEIPT	A message is accepted by Information Exchange for delivery to a receiver’s mailbox or valid distribution list after a successful session.
DELIVERY	Information Exchange generates a delivery acknowledgment when a destination user receives a message from the Information Exchange mailbox.
PURGE	Information Exchange generates a purge acknowledgment when a message is purged from the receiver's mailbox.

Although you can request acknowledgments by using the Send Message command, you must receive them from your Information Exchange mailbox by using the Receive Message command. A message number of 01000, 01001, or 01002 in MSGNUMBR (field 1) of the Information Exchange generated message text indicates an acknowledgment message.

For more information, see the discussion of Information Exchange acknowledgments in *Information Exchange Interface Programming Messages and Formats*.

Obtaining status information

You can use the following inquiry commands to determine how much data is currently available for retrieval. You can also use these commands to obtain session status information.

- Session inquiry

You can inquire any time during a session about the status of the session and receive an inquiry response. Information Exchange provides both a long- and short-format response to these commands. For more information, see “Session Inquiry command—SDIINQS” on page 99 and “Session Inquiry response” on page 100.

- Message inquiry

You can find out the amount of message data currently queued for your Information Exchange address by using the Message Inquiry command. For more information, see “Message Inquiry command—SDIINQM” on page 146 and “Message Inquiry response” on page 147.

Tracking messages and sessions

Audit trails are records maintained by Information Exchange that contain the status of messages that have been received or sent. Session trace records provide information about your Information Exchange session and the receive commands issued during the sessions. Both types of tracking provide useful problem determination assistance. You can work with audit trails and session traces by using Information Exchange Administration Services or Information Exchange interface commands.

Information Exchange has a defined minimum and maximum limit for record retention. Using Information Exchange Administration Services, you can set a storage time limit within these minimum and maximum values.

Message audit feature

Information Exchange maintains audit data about the status of your messages. Audit records are created for both the sender and each trading partner when the message is placed in a trading partner's mailbox. This information is updated when a message is received by a trading partner, purged, or canceled. You can browse and retrieve audit data to verify correspondence and transactions.

Two commands are provided for audit data retrieval:

- **Audit Browse command**

This command provides audit data that your application can browse. Information Exchange's response to this command is the Audit Browse response. This response contains the audit records that provide information on the status of your messages. For more information, see “Audit Browse command-SDIAUDB” on page 231.

- **Audit Retrieve command**

This command downloads the contents of your audit trail in the form of a message group into your mailbox. When the audit information is in your mailbox, you must issue the Receive Message command to receive it. Information Exchange's response to this command is the Audit Retrieve response. You can check this response to determine whether Information Exchange processed the command successfully. For more information, see “Audit Retrieve command—SDIAUDR” on page 237.

For more information on the data returned by the Audit command, see “Audit Record/Display format” on page 243. To determine the status of your messages, look at the value in ASTFLAG (field 5) of the audit record; see “ASTATFLG (field 5)” on page 246.

If a message has been purged, you can determine the reason by looking at the value in APURGFLG (field 7) of the audit record. Information Exchange purges a message as a result of a Cancel command by the sender or intended recipient, or when certain system-related events occur. See “APURGFLG (field 7)” on page 248.

Session trace feature

You can find information about your Information Exchange sessions by using the session trace feature. This feature provides detailed status information on your Information Exchange session. Two commands are provided for retrieval of session trace information:

■ **Session Trace Browse command**

This command provides session trace information that your application can browse. Information Exchange's response to this command is the Session Trace Browse response. This response contains session trace records that provide information on the status of your session.

For more information, see “Session Trace Browse command—SDITRLB” on page 104.

■ **Session Trace Retrieve command**

This command downloads session information you select in the form of a message group into your Information Exchange mailbox from *SYSTEM**EVENTS. When session trace information is in your mailbox, you must issue the Receive Message command to receive it. Information Exchange's response to this command is the Session Trace Retrieve response. You can check this response to determine whether Information Exchange processed your command successfully.

For more information, see “Session Trace Retrieve command—SDITRLR” on page 107.

Asynchronous background processing

Information Exchange performs the process that updates audit records asynchronously. For example, when you receive a message and commit its receipt either through using the checkpoint process or by ending the session, Information Exchange starts an asynchronous process to update the audit record for that message. If you look at the audit trail for the message immediately after receiving it, the message might still seem to be in the mailbox. Because of this asynchronous process, updated audit data and session trace information are not normally available until your next session with Information Exchange.

Information Exchange also starts an asynchronous process for the following reasons:

- To process the Cancel command
- To retrieve archives
- To process any Library command
- To process all administration mailbox commands
- To retrieve audits
- To retrieve session traces
- To determine the list of carbon copy recipients, if any
- To create a carbon copy or redirected message

Sending messages

This chapter describes how to send Information Exchange messages. This chapter also discusses several aids for message transfer, including using distribution lists, sending carbon copy messages, and validating transfer information.

Using the Send Message command

To send a message, use the Send Message (SDISNDM) command. This command contains both header information and the actual message text. Information Exchange considers the information you specify in the first 24 fields of the command to be the message header. The message header establishes how the system processes the message. The last field (field 25) of the command has a variable length and contains the actual message text.

The Send Message command includes information on the following:

- **Destination/address**

When you use the Send Message command, you must specify the address to which you want Information Exchange to send the message. You can specify the address of a single user or a list of users (distribution list).

- **Alias tables**

If you have several trading partners that you send messages to frequently, consider setting up aliases. Aliases are pseudonyms you can use instead of Information Exchange addresses. You can set up aliases that are short and easy to remember. For example, if a trading partner's address is a long alphanumeric code, set up an alias name for that trading partner that is the same as the trading partner's first name. Use the Send Message command to refer to the alias table that contains your trading partner's alias information (alias name and the complete Information Exchange address). Using aliases is easier and faster than typing a full Information Exchange address each time you send a message.

- **Type of normal header**

For message-segment continuation purposes, you can indicate whether you need to use a long or short header.

- **Message text length**

You can indicate the length of the message text immediately following the header information.

- **Message charges**

You can indicate who will pay for the message.

For a detailed explanation of the Send Message command, see “Send Message command—SDISNDM” on page 134.



NOTE: The Send Message command does not create a CDH, alias table, or distribution list; it only specifies their use.

Using the force selective receive search criteria

When sending a message, you can indicate whether or not the recipient of this message must specify selective receive search criteria for the message to be delivered. The SELRCV (field19) on the Send Message command is used for this purpose. A value of F in this field forces the receiver of the message to use one of the following selective receive search criteria:

- Sender’s address
- User message class
- Message key

The default is **blank**, indicating the message can be received without specifying the selective receive search criteria.

You should alert your trading partners when using this option so that they can specify search criteria and receive your message.

Message delivery classes

When you send a message, you can specify one of the following delivery classes:

- **Normal priority**

Information Exchange stores normal-priority messages until the addressee asks to receive them.

- **High priority**

Information Exchange stores high-priority messages until the addressee asks to receive them. Information Exchange usually delivers these messages before normal-priority messages. However, if a high-priority message is larger than a normal-priority message and your application has limited the amount of data it can receive from Information Exchange, the system might deliver the normal-priority message first.

- **Express priority**

Information Exchange does not accept express-priority messages unless the addressee has an appropriate Receive Message command in effect. Express priority is intended for situations where it is ineffective or unreasonable to send the message unless it can be delivered almost immediately. If Information Exchange cannot make an immediate delivery, it discards the message and informs the sender.

Addressing messages

You can address Information Exchange messages using the following:

- An explicit Information Exchange address that refers to a destination on your local Information Exchange system
- An explicit intersystem Information Exchange address that refers to a destination on a remote system
- An alias for a local Information Exchange address or intersystem Information Exchange address
- A distribution list

Information Exchange addresses

An Information Exchange address consists of an 8-byte account ID, followed by an 8-byte user ID (fields AAAAAAAAA and UUUUUUUU in Figure 4).

Intersystem Information Exchange addresses

Information Exchange's intersystem message facility lets you communicate with trading partners who are attached to different Information Exchange systems.

The format of an intersystem Information Exchange address is illustrated in Figure 4.

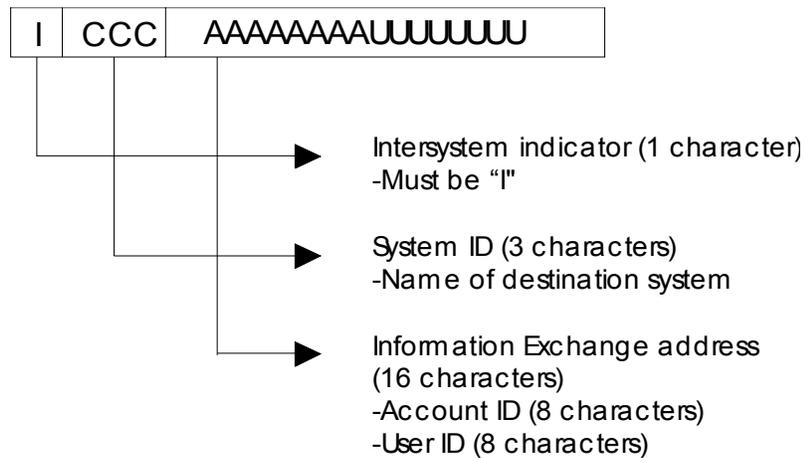


Figure 4. Format of an intersystem Information Exchange address

Specify the intersystem identifier (I) in DTBLTYP (field 23) and the system ID (CCC) in DTBLID (field 24) of the Send Message command. The system ID refers to a user on another system. Valid values for the system ID may include:

This ID	Indicates:
USA	USA, Canada, and Latin America
USQ	Australia and New Zealand
EUR	Europe, Middle East, Africa, and Hong Kong
JPN	Asian-Pacific countries, excluding those listed under USQ
BRA	Brazil
ADV	Used to reach trading partners connected to the Information Exchange Peer system, and trading partners connected to other EDI VANS
IMX	Used to reach trading partners connected to IBM Mail Exchange
MQS	Used to reach trading partners connected to IBM MQSeries Services
SRV	The EDI Server used to access the in-network translation feature of EDI Services, and to reach trading partners using Web data transfer, forms exchange, or Web application environment components of EDI Services

For the most current listing of system IDs, use Information Exchange Administration Services.



NOTE: In some commands, such as the Send Message command, the intersystem and system identifiers (I and CCC) are not contiguous with the fields that indicate the account ID and user ID (AAAAAAA and UUUUUUU in Figure 4).

Alias addressing

With Information Exchange, you can send information to another user by using an alias. Aliases are pseudonyms that you use instead of an Information Exchange address. When you use aliases to send or receive messages, Information Exchange converts them into Information Exchange addresses for its internal operation.

Information Exchange stores aliases in permanent tables. To create an alias table, use the Define Alias command. Aliases can be at any of the following levels:

- Individual user level (private)
- Account level
- Global level (can be addressed by any Information Exchange user on your local Information Exchange system)

You can use an alias to point to another alias or to a maximum of five other aliases. You can also determine the true Information Exchange address of any given alias using the Alias Inquiry command. For more information, see “Alias Inquiry command—SDIINQA” on page 176 and “Define Alias command—SDIDALS” on page 179.

Distribution lists

You can send the same message to more than one person at a time by making a list of users and sending the message to that list. You can create two types of distribution lists by using the List Define command: permanent and temporary.

Permanent distribution lists are permanently stored in Information Exchange. The advantage of this type of list is that many different users can access it. Permanent lists can be at an account level or at an individual user level.

Temporary distribution lists last only for the duration of your Information Exchange session. When your Information Exchange session ends, the system deletes your temporary list.

A list can include users on your local Information Exchange system and users on any number of other remote systems. You can express users in the form of Information Exchange addresses, intersystem Information Exchange addresses, aliases, or any combination of these.

When a command references a list, Information Exchange searches for the list, using the following hierarchy:

1. Individual's temporary distribution list created during the session
2. Individual's permanent distribution list
3. Account's permanent distribution list

If the system does not find the list name in any of these searches, the command fails.



NOTE: A list name consisting entirely of blanks is valid for temporary lists, but not for permanent lists.

See “List Define command—SDILSTD” on page 182 for the format of the List Define command. With the List Define command, you can:

- Define a new distribution list
- Replace a list
- Delete entries from a list
- Add entries to a list
- Erase an entire list

The same list can be referenced by both send and receive operations.

You can use the distribution list to receive messages originated by Information Exchange by placing the appropriate Information Exchange generated address such as *SYSTEM**ERRMSG*. For a complete list of system messages, see “System addresses” on page 58.

Verifying a distribution list

The List Define command does not verify the Information Exchange addresses on a distribution list. The first time a new list (or Information Exchange address) is referenced in a Send Message command, the Information Exchange addresses are tested for validity and authorization to receive messages from your Information Exchange address. Information Exchange does not validate the addresses on a distribution list in a Receive Message command.

To ensure that the Information Exchange addresses on a distribution list are valid and authorized to receive messages from your Information Exchange addresses, you must issue a List Verify command following the List Define command. For each unusual condition, the system places error messages from the List Verify command in your message queue.

To submit and verify a distribution list:

1. Define the list using the List Define command.
2. Submit a List Verify command.
3. Submit a Receive Message command to retrieve the results of the error messages from Information Exchange.

For more information, see “List Verify command—SDILSTV” on page 186.

Validating addresses, payment levels, and authorizations

To reduce transmission costs and ensure that your message can be delivered when communicating with unfamiliar trading partners, make certain that:

- The address to which you want to send the message is a valid address.

- You can use a particular payment level when sending the message to the intended destination.

You can tell Information Exchange to check the addresses, payment levels, and authorization levels before sending a message if you are on the same system as the intended destination. To do this, use the Probe command.

You can request an immediate response (synchronous), or you can have the response placed in your Information Exchange mailbox (asynchronous) so you can look at it later.

You can use the Probe command for distribution lists and for individual destinations. However, the Probe command only verifies that the distribution list exists. It does not verify the existence of individual entries within the distribution list.

If a verification succeeds, you do not receive a charge for the verification request. If a verification fails, you do receive a charge. For details, see “Probe command—SDIPROB” on page 127.



NOTE: You cannot use the Probe command to check authorization and payment levels for users who are not on your local Information Exchange system.

Organizing data for transmission efficiency

If possible, organize transactions that have the same destination or data type together. This usually increases efficiency and is less costly than sending each transaction individually.

Separating data by destination

Organize transactions that are to go to the same destination into a single Information Exchange message. If all data for a destination is available, you can submit it for delivery as a single logical message or as a message group.

The intended receiver of a message or message group must have the capacity to receive it in one Information Exchange session. Normally, the receiver must be able to store data until the entire message or message group has been received. At this point, it can be released for local processing.

Separating data by type

You can also separate transactions by data type or by processing requirements. If you identify the contents of the message in the MSGUCLS field of the Send Message command, the destination sorts the messages by type and processes them accordingly.

Suppose, for example, that you have four categories of data to send to another user:

- The first category consists of formatted records that need to be printed on special forms.
- The second category consists of administrative messages that are to be printed on plain paper.
- The third category is a complete set of file records to replace an existing inquiry file.
- The fourth category consists of update transactions to maintain a local master file.

You can send this data in a single logical message, but it may be simpler for both you and the receiver if the data is divided into messages, according to data type.

It may also be possible to respond more promptly to high-priority transactions, such as electronic memos, if they are clearly separated from batch data.

Making user data understandable

The meaningful exchange of information requires that all participants be able to understand the data they exchange.

Although Information Exchange provides machine protocol independence between message sender and receiver, it cannot provide data independence. Ultimately, the receiver must be able to interpret the message. You are required to establish a data protocol between communicating parties.

Self-defining data

You can make data records that you exchange with members of a business community self-defining. In other words, a record can contain data that identifies its data type. You can identify the fields within records by field name and content. There are two methods you can use to make data records self-defining.

One technique is to precede each data record or transaction element in the data stream with data type and length elements. You can then construct logical messages of one or more types of variable-length data records, and the recipient can start at the beginning of a message and locate the individual data units.

Another technique is to use structures with the basic format of "KEYWORD=value." For example:

```
NAME='BOB SMITH' ADDRESS='101 MAIN ST'  
CITY='ANYWHERE, USA' ZIP='11111'  
BALANCE=187.43
```

Here, various punctuation symbols permit the parsing of the message and the association of data element and content.

Data-element splitting

When you use message segments to transfer data, you can split data elements (fields) between these individual message segments. However, the recipient of a message must ensure that each data element is completely received before processing its content. In the above example, the string "ZIP='11111'" may have been split into "ZIP='11" and "1111" or "ZI" and "P='11111'."

Character-code sets

Each sender is responsible for ensuring that message data is properly encoded for the recipient.

Text compression considerations

Do not compress data in an Information Exchange command or in message headers. Use an asterisk (*) in command fields instead of blanks to prevent space compression within the command. The system replaces the asterisk with a blank before processing the command, except in the following fields that are used to specify a value you want returned in a subsequent request or command response:

- RESPTRN (field 6) of the Session Start command
- COMMCMD (field 12) of the Session Start command
- SENDRSPC (field 5) of the Session End command
- CMTRSPC (field 5) of the Commit command
- PACRSPC (field 5) of the Pace command
- QMSGTRID (field 10) of the Receive Message command
- IMSGTRID (field 11) of the Receive Message command
- SMSGTRID (field 12) of the Receive Message command
- INQSTRID (field 5) of the Session Inquiry command
- INQMTRID (field 5) of the Message Inquiry command
- RETRSPC (field 6) of the Audit Browse response
- RETRSPC (field 6) of the Archive Retrieve command
- RETRSPC (field 6) of the Audit Browse command
- RETRSPC (field 6) of the Audit Retrieve command
- QMSGTRID (field 7) of the Intermediate Destination Receive Message command
- RETRSPC (field 6) of the Message Queue Query command
- RETRSPC (field 6) of the Probe command

Send compressed data in the same way you send other types of data using the Send Message command.

Committing messages

If you specify checkpoint-level recovery, you can ask Information Exchange to periodically protect messages that have been submitted and establish a restart point in case of failure.

Messages that your interface submits to Information Exchange are not queued to their destinations until a recovery point is reached. In checkpoint-level recovery, a Commit exchange establishes a recovery point. This is accomplished by:

1. Sending a Commit command after one or more messages
2. Waiting for Information Exchange to respond to the Commit command
3. Reading and processing the Commit response that Information Exchange generates
4. Sending another message (or end data), a subsequent Commit command, or a Session End command.



NOTE: Do not issue a Commit command between message segments in a multiple-segment message. The previous Send Message command must have a *blank* or an *E* in MSGTXTOV (field 5), indicating the end of the message. If Information Exchange receives a Commit command, when the previous Send Message command specifies the message is to continue in subsequent segments (*S* or *C* in MSGTXTOV—field), the session abends. However, you can issue a Commit command between the completed messages of a message group.

Considerations for committing messages

Information Exchange uses a two-stage commit process, in which the messages are not actually committed or forwarded to their destination until after Information Exchange receives a Send Message command, another Commit command, or a Session End command. The system does not deliver messages until you do the following:

1. Send the message to Information Exchange with a Send Message command.
2. Send a Commit command.
3. Receive a Commit response from Information Exchange.
4. Send one of the following:
 - a. A Send Message command for a message or ENDDATA to Information Exchange
 - b. A Session End command to Information Exchange for a normal session end
 - c. Another Commit command

Commit response

Information Exchange sends a Commit response to your Information Exchange address in response to a Commit command from that Information Exchange address. For more information, see “Commit command response” on page 124.

For the format of the Commit response sent by Information Exchange, see “Commit command response” on page 124. See “Controlling the Information Exchange session” on page 11, for information about processing these messages.

Requesting a pause in transmission—Pace command

With Information Exchange, you can send a Pace Command to cause a transmission synchronization pause after any completed message; the session does not need to be using checkpoint-level recovery. Information Exchange responds to the Pace command with a Pace response, indicating that all messages up to, and including, the Pace command have been processed.

The Pace command works like the Commit command, but it does not cause the messages to be protected as the Commit command does. The Pace command does not provide a restart-and-recovery point. For the format of the Pace command, see “Pace command—SDIPACE” on page 125. To use the Pace command:

1. Send a Pace command after one or more messages.
2. Wait for Information Exchange to respond to the Pace command.
3. Read and process the Pace response that Information Exchange generates.



NOTE: Do not send a Pace command between message segments of a multiple-segment message. (The previous Send Message command must have a *blank* or an *E* in MSGTXTOV.) However, you can issue a Pace between the completed messages of a message group.

Pace response

Information Exchange sends a Pace response to your Information Exchange address in response to a Pace command from that Information Exchange address.

For the format of the Pace response, see “Pace command response” on page 126.

Receiving messages

This chapter describes how to receive messages that are queued to an Information Exchange address. It also describes what you can expect while receiving messages. Included are:

- Messages from other Information Exchange users
- Commit requests and responses
- Pace requests and responses
- Messages generated by Information Exchange

Using the Receive Message command

To receive messages, use the Receive Message command (SDIRCV). Before you can use the Receive Message command, you must start an Information Exchange session by using the Session Start command. The Receive Message command uses the values you specify in several Session Start command parameters to process your message. For more information on the effects of the Session Start command on receiving messages, see “How the Session Start command affects receiving messages” on page 13.

The Receive Message command lets you specify from whom you want to receive messages:

- All users
- A single user
- A list of users

You can also specify how much you want to receive:

- All messages
- Selected messages

You can limit the amount of message segments you want to receive for a message group by specifying a value in MAXMSGNO (field 14) of the Receive Message command. In addition, you can specify that the command remains active after the messages you selected are received. This mode is called continuous receive. Information Exchange forwards additional messages as they arrive. For more information, see “Receive Message command—SDIRCVM” on page 161.

Receiving messages sent using force selective receive search criteria

Using the SELRCV (field 19) on the Send Message command, the sender of a message can force the receiver to specify search criteria for the message to be delivered.

You can determine if a message has been sent with this option by viewing the mailbox in Information Exchange Administration Services. Messages sent with this option will display an "F" in the Rcv ind field on the list of inbound messages panel. You can also use the Information Exchange Message Queue Query command to examine inbound messages. Issue the command with expansion level 2. An "F" is returned in RCVIND (field 12) of the Message Queue entry response if force selective receive search criteria was used when the message was sent.

Format of received messages

The format of the messages you receive from Information Exchange is variable in length and has two parts:

- Message header

The message header serves to identify the type and source of the message.

- Message text

The message text is the text as sent from the submitting system, with any specified transformations applied to it.

Using the Commit request

If you specify checkpoint-level recovery in your Session Start command, you receive Commit requests from Information Exchange during message retrieval. The frequency of these Commit requests is determined by the Session Start command field MAXMSGCP (field 11). See “Controlling the Information Exchange session” on page 11, for information about this frequency.

You should not anticipate the arrival of a Commit request. Information Exchange always sends a Commit request following an ENDDATA message. For the format of the Commit request that Information Exchange sends, see “Commit request” on page 142.



NOTE: Information Exchange does not send the Commit request between segments of a multiple-segment message. However, it may send a Commit request between the completed messages of a message group.

Responding to a Commit request

The Commit response must be sent to Information Exchange when you receive a Commit request from Information Exchange. Information Exchange will not send any more messages to an Information Exchange address following a Commit request until a Commit response is received by Information Exchange.

For the format of the Commit response, see “Commit command response” on page 124. See “Controlling the Information Exchange session” on page 11, for information about processing these messages.

Understanding the Pace request

If the session does not use checkpoint-level recovery, your application receives Pace requests from Information Exchange during message retrieval. The MAXMSGCP (field 11) of the Session Start command determines the frequency of the Pace request. See “How the Session Start command affects receiving messages” on page 13 for more information about this frequency. Information Exchange always sends a Pace request following an ENDDATA message.

When you receive a Pace request from Information Exchange, you must answer with a Pace response; Information Exchange cannot send you any more messages until you have done this. For the format of the Pace request that Information Exchange sends to your application, see “Pace request” on page 157.



NOTE: Information Exchange does not send the Pace request between segments of a multiple-segment message. However, it may send a Pace request between the completed messages of a message group.

Responding to a Pace request

The Pace response must be sent to Information Exchange when you receive a Pace request from Information Exchange. Information Exchange will not send any more messages to an Information Exchange address following a Pace request until it receives a Pace response.

For the format of the Pace response, see “Pace command response” on page 126.

Messages generated by Information Exchange

Information Exchange generates messages to be sent to the user's Information Exchange mailbox. See *Information Exchange Messages and Formats*. These may be information messages, such as acknowledgments, or they may be error messages. Information Exchange sends them to the mailbox in a format similar to those for data messages from other users. The header contains:

- The value from SMSGRID of the Receive Message command in its COMMAND field
- The value *SYSTEM* in its DESTACCT field
- A value in its DESTUID field which indicates the type of message. See “System addresses” in this chapter for a list of possible values.

System addresses

The message, as it resides in the mailbox, displays the sender as account = *SYSTEM*, and the user ID as follows:

This user ID:	Indicates:
*ADMIN**	Message from commands submitted to the Administration mailbox
ADMINGET	Output from Get requests to the Administration mailbox
AUDITS	Audit records
ERRMSG	Error messages or message acknowledgments
EVENTS	Session trace records
LSTRSP	Output from list verify
PRBRSP	Asynchronous probe response
TSTMSG	Test messages

An example is when the user retrieves an audit from Information Exchange. Information Exchange places the audit file in the user's mailbox with the sender's address of *SYSTEM* *AUDITS*. Users must issue the Receive Message command to move the file from the Information Exchange mailbox to their system environment.

The following user IDs are ONLY found on a user's invoice.

This user ID:	Is generated when:
ADTBRS	Browsing audits on Information Exchange Administration Services
CALLOUT	An event notification for a call-out user is sent to the user.
DIALOUT	An event notification for a dial-out user is sent to the user.
*LIBBRWS	Browsing library members on Information Exchange Administration Services
LIBSTOR*	Storing members in a library
MSGBRS	Viewing selected messages in the mailbox on Information Exchange Administration Services



NOTE: The network charges you for receiving messages from the above addresses, except for the error message address.

Archiving messages

One of Information Exchange's features is the ability to archive messages when they are received. This allows you to retrieve a previously received message at a later time.

When a user profile is created, there are two fields that affect archiving:

- Forced archiving
- Number of archive days

The number of days an archive is kept is determined by the value specified in the "number of archive days" field in the user's profile. The system prevents a user from setting Forced archiving to *Y* and setting the number of archive days to zero.

Here are ways to cause a message to be archived:

1. "Forced archiving" is set to *Y*:
 - When a user receives a message, it will be archived.
 - Where an archive group reference ID is specified in the Receive Message command, the message will be archived using the specified archive group reference ID.
 - Where an archive group reference ID is not specified, the session access key will be used as the archive group reference ID when archiving the message.
2. If "Forced archiving" is set to *N*:
 - A message will be archived only if an archive group reference ID is specified when the Receive Message command is issued. In this case, the message is archived only if the number of archive days is greater than zero.

Creating groups of archive messages

One way to group archive messages is by the way the messages are received. Using an archive reference can narrow the criteria used to find archived messages. Selecting an archive reference, other than the receive date, is recommended. Some options are:

- If MSGCLASS is used in the selection criteria of the Receive Message command, you may want to consider using MSGCLASS as the archive reference.
- If you select messages to be received by sender's ID, you may want to use an archive reference that relates to the sender's ID. This way you can group all messages received from the same sender.
- If you receive messages based on the time the messages were sent, you may want to group messages by date sent.

Retrieving messages from archive is easier when the messages are grouped together by similar criteria.



NOTE: You should avoid using the same archive group reference ID for all your receives. This minimizes searching through multiple panels in Information Exchange Administration Services at archive message retrieval time.

Archiving carbon copy messages

You can request archiving of messages when setting up receive-side relationships for carbon copy messages. For further information on this procedure, see the *Information Exchange Administration Services User's Guide*.

Library services

A library is a place to store data for an extended period of time. For example, you can store:

- Product catalog information
- Engineering specifications
- Problem descriptions
- Programs

The account defined when the library is created is the owner of that library. The owner can:

- Specify the name of the library
- Identify the owning account ID and user ID
- Provide a description of the library
- Indicate whether the library is searchable
- Establish the charges structure
- Grant users read- and write-level access
- Grant network personnel access to the library

The owner authorizes users to access and work with the library. This access is granted separately for read and write authority. Read authority lets the user access, browse, and retrieve the data contained in the library; however, this user cannot make changes to the library information. Write authority lets the user access, browse, retrieve, and add and delete library members.

Units of data contained in the library are called members. A member can be a message or file that you want many different users to access. A library can consist of one or more members and is permanent until you explicitly delete it. To place a member in the library, use the Send Message command. For more information, see “Adding members to the library” on page 62.

Library functions

The following sections describe the library commands. Command formats are found in “Library commands” on page 189.

Defining a library

Use the Define Library command (SDILDEF) to define a library. This command defines the characteristics of the library, such as whether the library is searchable and who has read or write authority. The Define Library command returns the Define Library response. This response indicates the command processing through the value specified in RETCODE (field 6).

Adding members to the library

To add a member to the library, you must:

1. Have update access to the library
2. Use the Send Message command and specify the destination `"*SYSTEM**LIBRARY"`
3. Specify the following CDH fields:
 - **Library owning account (LIBACCT)**
Indicate in this field the owning account name of the library that will contain the data. The default is your account.
 - **Library name (LIBNAME)**
Indicate in this field the name of the library that will contain the data.
 - **Library member (LIBMEMB)**
Indicate in this field the name of the library member that will contain the data.
 - **Library replacement option (LIBREPL)**
Indicate in this field whether the member you are adding is replacing an existing member with the same name. If you do not specify the replace option and a member with the same name exists, Information Exchange generates an error message.

Deleting members from the library

To delete a member from the library, use the Delete Library Member command (SDIDLDM). Specify the name of the member you want to delete in MEMBER (field 9). The Delete Library Member command returns the Delete Library Member response containing information on command processing.

Deleting a library

To delete a specific library, use the Delete Library command (SDILDEL). Use the force indicator if you want to delete a library and all its members by specifying *F* in OPTIONS (field 9).

The Delete Library command returns the Delete Library response containing information on command processing.

Retrieving a library member

To place a copy of a library member in your mailbox, use the Retrieve Library Member command (SDILRTV). You can direct library members to your mailbox by using standard Information Exchange destination specifications. Information Exchange places the member in your mailbox as a message group. You must receive this message group from your mailbox by using the Receive Message command. Charges are based on receive-side payment levels. Refer to *Information Exchange Charges Reference* for message charge resolution.

Updating library members

To update a library member, do the following:

1. Retrieve the library member using the Retrieve Library Member command.
Information Exchange places the library member in your mailbox as a message group.
2. Receive the library member from your mailbox using the Receive Message command.
3. Update the library text.
4. Add the library member back to the library by using the Send Message command and specifying the destination as *SYSTEM**LIBRARY*. And then specify the following CDH fields:
 - **Library owning account (LIBACCT)**
Indicate in this field the owning account name of the library where you want to store the member.
 - **Library name (LIBAME)**
Indicate in this field the name of the library where you want to store the member.
 - **Library member name (LIBMEMB)**
Indicate in this field the name of the library member.
 - **Library replacement option (LIBREPL)**
Specify the replace option in this field. If you do not specify this option, Information Exchange generates an error message.

Browsing library members

To retrieve a partial text of a library member and browse its contents, use the Browse Library Member command (SDILBRW). This command returns a Browse Library Member response that contains a partial text of the library member you want to browse.

The Browse Library Member response also contains a location indicator. Copy the value from this indicator, which is specified in NEXTLOC (field 8), into LOCATOR (field 11) of your subsequent Browse Library Member command when you want Information Exchange to return to you the next piece of text in the library member from the last breakpoint.

Listing libraries

Use the List Libraries command (SDILLIST) to list account libraries or libraries for which you have read or write access. This command returns a List Libraries response that provides the library listing and statistics that relate to the library.

The List Libraries response also contains a location indicator. Copy the value from this indicator, which is specified in NEXTLOC (field 8), into the forward location indicator (FORWLOCT—field 10) or the backward location indicator (BACKLOCT—field 11) of your subsequent List Libraries command when you want the system to continue listing the libraries from the last breakpoint. Because of the maximum message size specification you indicated in MAXMSGSZ (field 10) of the Session Start command, Information Exchange may not have been able to send you a complete listing of the libraries. To determine whether the library list continues, look at the value in CONTINUE (field 7) of the List Libraries response. A value other than 0 indicates there are more libraries to follow. Issue subsequent List Libraries commands in the above manner, until the value in CONTINUE of the List Libraries response is 0.

For more information on the structure of the library items, see “Library Entry item structure” on page 207.

Listing library members

To list library members and statistics, use the List Library Members command (SDILLIST). This command returns a List Library members response that provides information on command processing and a list of library members.

The List Library Members response also contains a location indicator. Copy the value from this indicator, which is specified in NEXTLOC (field 8), into LOCATOR (field 9) of your subsequent List Library Members command when you want the system to continue listing the members from the last breakpoint. Because of the maximum message size specification you indicated in MAXMSGSZ (field 10) of the Session Start command, Information Exchange may not have been able to send you a complete listing of the members. To determine whether the member list continues, look at the value in CONTINUE (field 7) of the List Library Members response. A value other than 0 indicates there are more members to follow. Issue subsequent List Library Members commands in the above manner until the value in CONTINUE of the List Library Members response is 0.

For information on the structure of the member list items, see “Library Member List entry” on page 213.

Searching the library

There are two types of libraries: searchable and non searchable. If you have read access to a searchable library, you can search a library for members that meet your search criteria by using the Search Library command (SDILSCH). The Search Library command returns a Search Library response containing an initial set of members that match your search string. The size of that set and the number of members it contains depends on the value you specified in MAXMSGSZ (field 10) of the Session Start command. This field indicates the maximum message size your system can receive. The Search Library command may not be able to return all the members that meet your search criteria, because of maximum message size specifications. To determine whether there are additional members, look at the value in NEXTLOC (field 7). A series of Fs in NEXTLOC indicates there are no more members to follow. A value other than a series of Fs indicates there is another set of members that meet your search criteria.

Issue a Library Title command to retrieve additional members. Copy the location indicator value in NEXTLOC of the Search Library response into LOCATOR (field 7) of the Library Title command to let the system locate the next set of members from the last breakpoint. The Library Title command returns the Library Title response, which contains a list of the additional members and information on command processing. Look at the value in NEXTLOC (field 7) to determine whether there are still more members in the list to be retrieved. Continue issuing subsequent Library Title commands in the above manner, until NEXTLOC contains a series of *F*s.

Handling errors

When Information Exchange encounters an error in your command input, it usually creates an error message and places it in your queue of messages. However, certain errors make it impossible for Information Exchange to place an error message on your message queue. For example, when an input message header arrives at Information Exchange with an invalid Information Exchange address, Information Exchange determines that there can be no message queue. In such cases, Information Exchange sends back an error response message to the originator of the invalid input. Information Exchange cannot process the input beyond determining that it is invalid.

Information Exchange sends this type of error response message when the following occurs:

- The command name provided in the input is invalid.
- The input command is too short to contain all necessary fields.
- An Information Exchange session was not established for the Information Exchange address.
- The Information Exchange session access key is invalid.

The commands that normally receive responses, such as Session Start, Session End, Commit, Pace, Session Inquiry, and Message Inquiry, may also receive this error response message instead of the normal response.

Error response message format

The format of the error response message is shown in Figure 5.

#	Col	Size	Name	Description
1	1	8	COMMAND	Command name
2	9	48	TEXT	Constant message text
3	57	5	ERRCODE	Error message number
4	62	64	ERRTXT	Erroneous input data

Figure 5. Error response message. Information Exchange sends this 125-character message when it cannot process your input because that input contains serious errors.

COMMAND (field 1)

This field identifies this message as an error response. It contains the character string "SDIERRbb".

TEXT (field 2)

This field contains the constant text string "THE FOLLOWING COMMAND WAS NOT EXECUTED BECAUSE: b".

ERRCODE (field 3)

This field contains an error code identifying the cause of the error.

This value:	Indicates:
1	The first 8 bytes of the input did not contain a valid Information Exchange command name, or you specified a command you are not authorized to use.
2	The input is too short to process; it did not contain all required fields.
4	No session was previously established for the Information Exchange address.
8	The session access key is invalid.
16	The expansion level is invalid on a varying length command. The default cannot be determined.
32	Serious errors in command fields are found. Issue a Receive Messages command to determine which fields are invalid.
64	The input was not processed. Try again. If the problem persists, contact the Customer Care Help Desk.

ERRTXT (field 4)

This field contains the first 64 characters of the erroneous command to which this message is responding.

Common data header

Information Exchange applications can use a common data header (CDH) to communicate detailed information about files and messages to other systems. The CDH provides details (such as file name and carriage-return and line-feed options) that enable the receiving application to reconstruct a received message to its original format. It also makes more information available to the recipient of the message group.

If the CDH is present, it must be the text of the first message of a message group. Whenever a CDH is used, MSGNCLS of the long header is set to **S** for normal mode, or to **R** for test mode. For more information, see “Common data header (CDH)” on page 23.

CDH entry fields

This section describes the CDH entry fields. If possible, program your interface to support all fields. However, when you send a CDH, you only need to specify those fields necessary for your particular transmission.

The CDH always begins with a 2-byte length field (HLENGTH) that describes the length of the entire header, including the 2-byte length field. The CDH is comprised of variable-length entries that are order-independent, except for HLENGTH. The following table illustrates only the maximum length of each entry data. The actual length of each CDH field is 3 bytes more than illustrated in the table. Each entry contains a 2-byte length field and 1-byte ID (hex values), in addition to the variable length of the entry data.

ID	Name	Size	Description
1	FILENAME	54	Sender's file identifier
2	LOCATION	65	Sender's file location
3	RECORDF	4	Record format
4	RLENGTH	2	Record length

ID	Name	Size	Description
5	CTYPE	1	Data type
6	TTYPE	8	Translate type
7	DELIMIT	1	Record delimiter
8	DFORMAT	1	Data format
9	DESCRIBE	79	Description of message
A	UNIQUEID	8	Unique ID
B	CODEPAGE	8	Sending code page
C	SYSTYPE	1	Sending system type
D	VERSION	1	Sending system version
E	FILEDATE	6	File date
F	FILETIME	6	File time
10	CRLFEOF	3	CRLF and end-of-file (EOF) characters
11	Reserved	1	Reserved
12	EDISQUAL	4	EDI sender ID qualifier
13	EDISENDER	35	EDI sender ID
14	EDIRQUAL	4	EDI receiver ID qualifier
15	EDIRECVR	35	EDI receiver ID
16	EDICNTLN	14	EDI interchange control number
17	ORGLEN	4	Length of data
18	LIBACCT	8	Library owning account
19	LIBNAME	8	Library name
1A	LIBMEMB	8	Library member
1B	LIBREPL	1	Library member replacement
1C	COMPSOFT	10	Software used for compression
1D	COMPSVER	5	Version of software used for compression
1E	COMPDLEN	10	Length of the data before compression
1F	COMFNAME	54	Compressed file name

Figure 6. CDH entry fields

FILENAME (ID=X'01')

If the message is sent from a file, this is the system-dependent file name on the originating system. If the message is not sent from a file, this entry does not exist in the CDH. An interface uses this field to generate a receive file name that matches the send file name when sending data from one personal computer (PC) to another. When sending data between different systems, this field is primarily user documentation.

Use 1 to 54 alphanumeric characters.

LOCATION (ID=X'02')

This is the location of the file on the originating system. The value this field contains depends on the type of system you are using.

For this system:	Use a value that indicates:
MVS	The device type and volume serial number
VM	The minidisk label
AIX	The path
PC	The drive or directory

Use 1 to 65 alphanumeric characters. The default is blank.

RECORDF (ID=X'03')

This is the record format of the file on the sending system. (Use 1 to 4 alphanumeric characters.)

RLENGTH (ID=X'04')

This is the record length of the file on the sending system. (Use a 2-byte binary value.)

CTYPE (ID=X'05')

This is the type of data in the message. The data can be either EBCDIC or binary. For non-PC devices, this field is primarily documentation. For PC devices, this field controls translation. The PC, which is an ASCII device, must translate non-binary data to EBCDIC before it sends the data to Information Exchange. The PC must also translate non-binary data to ASCII after receiving it from Information Exchange.

This value:	Indicates the message contains:
1	EBCDIC data
2	Binary data

Use a 1-byte binary value.

TTYTYPE (ID=X'06')

Use this field to specify the translation table you want your system to use when it translates data between ASCII and EBCDIC. The translation table must be contained in a 256-byte file with an extension of XLT. For example, if the entry contained "3270" on receiving a file, the system loads the translate table from the file name 3270.XLT before translating the received data.

Use 1 to 8 alphanumeric characters. If this entry does not exist in the CDH, the system uses the default of **IE** when translating data.

DELIMIT (ID=X'07')

When you send through Information Exchange, the system does not retain its record format. This field contains the record delimiters the sending interface provides. When you receive data, your interface reconstructs the record format by using the information contained in this field.

This value:	Indicates that you:
1	Delimit records with CRLF characters
2	Precede records with a 2-byte length
3	Use EDI delimiters when reconstructing records
4	Do not delimit records. This is the default for data that is not in EDI format

Use a 1-byte binary value.

DFORMAT (ID=X'08')

Use this field to describe the format of data contained in the message. You can specify whether the format is non-EDI, EDI, X12, UCS, or UN/TDI.

This value:	Indicates that the data format is:
1	X12
2	UCS
3	EDIFACT
4	UN/TDI
5	Non-EDI

Use a 1-byte binary value.

DESCRIBE (ID=X'09')

Use this field to contain a free-format description of the message.

Use 1 to 79 alphanumeric characters.

UNIQUEID (ID=X'0A')

This field is the random ID assigned by the sending interface. This field is for documentation purposes and enables you to identify the message.

Use 1 to 8 alphanumeric characters.

CODEPAGE (ID=X'0B')

The character representation of symbols, such as \$ and #, may have different meanings in different countries. This field provides the receiver with the intended meaning of the character. This is called code page information.

Use 1 to 8 alphanumeric characters.

SYSTYPE (ID=X'0C')

This field indicates the type of system that is sending the data. This field enables the receiving interface to process received data differently based on the sending system type.

This code:	Means:
01	Unknown System Type
10	expEDite/PC
11	Expedite Base/2
12	Expedite Base/AIX
14	Expedite Base for SCO UNIX
15	Expedite Base/DOS
16	Expedite Base for SCO XENIX
17	Expedite Base for Windows
19	Expedite for Windows
20	expEDite/MVS Host
21	Expedite Base/MVS
22	TCP/IP FTP Gateway
30	IBM Mail Exchange
31	Expedite Base/VM
33	X.400 Gateway
40	Expedite/Direct
44	EDI VAN Interconnect
60	IBM MQSeries Services
61	EDI Server
71	Expedite Base/400
80	Expedite/CICS

This code:	Means:
90	Information Exchange Administration Services
91	Expedite/Async

Use a 1-byte binary value.

VERSION (ID=X'0D')

This field indicates the software version of the system sending the data.

Use a 1-byte binary value.

FILEDATE (ID=X'0E')

This field contains the date the file was last modified on the sending system. It is in the format YYMMDD.

Use 6 numeric characters.

FILETIME (ID=X'0F')

This field contains the time the file was last modified on the sending system. It is in the format HHMMSS.

Use 6 numeric characters.

CRLFEOF (ID=X'10')

This field contains the carriage-return line feed (CRLF) and end-of-file (EOF) characters you use when you want the file to be delimited by CRLF. Use this field when sending data to a PC that receives with a translate table other than the standard Information Exchange translate table.

RESERVED (ID=X'11')

This field is reserved; leave it blank.

EDISQUAL (ID=X'12')

This field contains the ID qualifier of the EDI sender. For X12 data, this field contains a 2-byte qualifier. For EDIFACT data, it contains a 4-byte qualifier. This field contains the entire qualifier, including trailing blanks. This field is not used for UCS and UN/TDI data, because they do not have an ID qualifier.

Use 1 to 4 alphanumeric characters.

EDISENDER (ID=X'13')

This field contains the ID of the EDI sender. It contains the entire ID, including trailing blanks.

Use 1 to 35 alphanumeric characters.

EDIRQUAL (ID=X'14')

This field contains the ID qualifier of the EDI receiver. For X12 data, this field contains a 2-byte qualifier. For EDIFACT data, it contains a 4-byte qualifier. This field contains the entire qualifier, including trailing blanks. This field is not used for UCS and UN/TDI data, because they do not have an ID qualifier.

Use 1 to 4 alphanumeric characters.

EDIRQUAL (ID=X'14')

This field contains the ID qualifier of the EDI receiver. For X12 data, this field contains a 2-byte qualifier. For EDIFACT data, it contains a 4-byte qualifier. This field contains the entire qualifier, including trailing blanks. This field is not used for UCS and UN/TDI data, because they do not have an ID qualifier.

EDIRECVR (ID=X'15')

This field contains the ID of the EDI receiver. It contains the entire ID, including trailing blanks.

Use 1 to 35 alphanumeric characters.

EDICNTLN (ID=X'16')

This field contains the interchange control number for EDI data.

Use 1 to 14 alphanumeric characters.

ORGLEN (ID=X'17')

This field contains the original length of the data before the addition of any record delimiters (on an MVS system). A value of zero indicates the length is unknown.

Use a 4-byte binary value.

LIBACCT (ID=X'18')

This field contains the owning account of the library that contains the data.

Use 1 to 8 alphanumeric characters. The default is the sending user's account.

LIBNAME (ID=X'19')

This field contains the name of the library that contains the data. You must supply this field when you want to reference a library. If you supply LIBNAME, you must also supply LIBMEMB.

Use 1 to 8 alphanumeric characters.

LIBMEMB (ID=X'1A')

This field contains the name of the library member that contains the data. You must supply this field when you want to reference a library. If you supply LIBMEMB, you must also supply LIBNAME.

Use 1 to 8 alphanumeric characters.

LIBREPL (ID=X'1B')

This field states whether the library member should be replaced. If you do not specify LIBREPL and the member exists, Information Exchange generates an error. A value of Y tells the system to replace the member. A blank or a value of any character other than Y tells the system not to replace the member. This is the default.

Use 1 alphanumeric character.

COMPSOFT (ID=X'1C')

This field indicates the software used to compress the data.

Use 1 to 10 alphanumeric characters.

COMPSVER (ID=X'1D')

This field indicates the version of the software used to compress the data.

Use 1 to 5 alphanumeric characters.

COMPDLEN (ID=X'1E')

This field indicates the length of the data after compression. This field is in EBCDIC numeric character representation. For example, X'F1F2' represents a length of 12.

Use 1 to 10 numeric characters.

COMFNAME (ID=X'1F')

This field indicates the name of the file where the compressed data has been stored.

Use 1 to 54 alphanumeric characters.

Cluster mailbox

With the cluster mailbox feature of Information Exchange, you can define a set of mailboxes that appear as one mailbox to your trading partners. This feature allows you to define additional mailboxes for use in sending and receiving data through Information Exchange without requiring your trading partners to change address information. This feature is most beneficial if you send and receive large amounts of data from many trading partners while operating within a restricted window for completing message transport.

Cluster user terminology and naming convention

A *cluster user* is a user who is known to trading partners by a generic name referred to as the *cluster parent user ID* (parent user ID). A cluster user can have from 2 to 99 mailboxes associated with that parent user ID. The individual mailboxes are referred to as *cluster child user IDs* (child user ID). The user ID of the parent is limited to 5 characters. The next two characters are used to assign a sequence number suffix to the parent user ID. IDs with this two-character suffix then become the child user IDs.

For a parent user ID named:	The child user ID names are:
ACCOUNT.USER	ACCOUNT.USER01
	ACCOUNT.USER02

In the example, ACCOUNT.USER is the parent user ID. ACCOUNT.USER01 and ACCOUNT.USER02 are the child user IDs.

Cluster user registration

Cluster users are defined as a result of a registration order.

The following options are available:

- Convert an existing user to a cluster user
- Add a user as a cluster user
- Convert a cluster user to a non-cluster user
- Change the number of mailboxes in a cluster

Cluster users can only be defined as general users to Information Exchange, and do not have access to Information Exchange Administration Services. Another non-cluster user must be authorized to perform service administrator activities for the cluster user.

Special considerations for cluster users

The cluster user starts sessions with Information Exchange using the child user IDs. The parent user ID is not allowed to establish sessions with Information Exchange and is treated as an invalid user ID.

As a cluster user, you must manage the sessions with Information Exchange for each of the child user IDs. This allows simultaneous processing to be achieved.

When sending messages to Information Exchange, you can distribute the messages to be sent among the available child user IDs. This can result in faster completion of message transport. If acknowledgments are requested when sending a message, they are placed in the mailbox of the child user ID issuing the Send Message command.

If multiple child user IDs are used to send messages to a single trading partner, the messages arrive at that trading partner's mailbox in an indeterminate sequence even if the "deliver messages in sequence" option is set in the profiles. If you need delivery to be sequential, you should not send to a single trading partner from multiple child user IDs.

When receiving messages, it is important that the Receive Message command be issued for all the available child user ID mailboxes. This is because a message destined to a parent user ID is placed in one of the child user ID mailboxes based upon a hash of the sender's ID so that all messages to a cluster user from a single trading partner go to a single child user ID. Issuing the Receive Message command for all child user ID mailboxes ensures all the mail items available to be received by the user are delivered.



NOTE: As a cluster owner, do not make any assumptions concerning the placement of mail into a specific child user ID mailbox by Information Exchange. The algorithm used to determine the placement of the message may be changed.

Information Exchange associates several objects with a particular type of user ID. The particular type of user ID, parent or child, associated with each of these objects is shown in the following:

Object name:	User ID type:
User profile	Cluster parent
User groups	Cluster parent
Alias tables	Cluster parent
Permanent private distribution list	Cluster parent
Temporary distribution list	Cluster child
User mailbox	Cluster child
Archived messages	Cluster parent
Message audits	Cluster parent
Session traces	Cluster child
Trading partner list	Cluster parent
Libraries	Cluster parent
Information Exchange user's X400 user profile	Cluster parent
X400 trading partner attributes	Cluster parent
Event profile	Cluster parent or child
Message arrival definitions	Cluster parent
Schedule event definitions	Cluster parent
Event traces	Cluster child
User passwords	Cluster child
User sessions	Cluster child
Billing records	Cluster parent
Carbon copy relationships	Cluster parent

Figure 7. Objects and associated User ID types

Handling commands

Warning and Error messages resulting from commands issued to Information Exchange are placed in the mailbox of the child user ID issuing the command.

Asynchronous responses to the following commands issued by child user IDs are placed in the mailbox of the child user ID issuing the command:

- SDILSTV—List Verify Command
- SDIPROB—Probe Command
- SDIAUDR—Audit Retrieve Command
- SDITRLR—Session/Receive Trace command

- SDILRTV—Retrieve Library Member Command
- SDIRCDH—Retrieve Library Member CDH
- SDILTST—Load Test Messages

Archiving messages

The message archive is owned by the parent user ID. When any of the cluster child user IDs receive messages requesting archive, the messages are placed inside a group under the parent user ID's archive. The archive retrieve command can be issued by any of the available child user IDs. The archive group is placed into the mailbox of the child user ID issuing the command. You must be careful not to request the same archive group to be retrieved into multiple child user ID mailboxes. If this is requested, the archive group is retrieved into all the requested mailboxes and could result in duplicate processing of a message by the cluster owner.

Audit Trails

When processing the following audit commands, you can request the audits for all child user IDs by specifying the cluster parent user ID in the ALTUSRID field of the command.

- SDIAUDR—Audit Retrieve command
- SDIAUDB—Audit Browse command

If the ALTUSRID field is left blank, only the records for the child user ID requesting the audit will be returned. The record/display format of the audit returned always specifies the parent user ID as the owner.

Session/Receive traces

You can request the traces for all child user IDs when processing the following session/receive trace commands:

- SDITRLR—Session/Receive Trace Retrieve
- SDITRLB—Session/Receive Trace Browse

Specify the parent user ID in the ALTUSRID field of the command. If the ALTUSRID field is left blank, only the records for the child user ID requesting the trace are returned. The record format of the trace returned specifies the child user ID associated with the trace.

Changing the number of cluster mailboxes

Active or incomplete sessions for any existing child user IDs must be completed or closed before a request to change the number of cluster mailboxes can be handled.

When the number of mailboxes changes, mail items may be moved to a different reset child mailbox. If the number of cluster child mailboxes is decreased, acknowledgments destined to a child user ID no longer available are placed in the first child user ID mailbox. If a message is sent to another Information Exchange system that does not support the cluster mailbox function, the acknowledgment is placed in the first child user ID mailbox.

Information Exchange commands and responses

This chapter describes Information Exchange commands and any responses these commands generate. This chapter also includes a brief description of the sequence of commands and tells you how to expand existing commands to accommodate additional options and future releases of Information Exchange.

If you are a cluster mailbox user, see “Special considerations for cluster users” on page 78.

Commands that cause responses

The following commands cause Information Exchange to send a response to the submitting system, indicating the results of the command processing. You do not need to issue a Receive Message command to receive a response from these commands.

- Alias Inquiry (SDIINQA)
- Archive Retrieve (SDIARTV)
- Audit Browse (SDIAUDB)
- Audit Retrieve (SDIAUDR)
- Browse Library Member (SDILBRW)
- Commit (SDICMIT)
- Define Library (SDILDEF)
- Delete Library (SDILDEL)
- Delete Library Member (SDILDLM)
- Library Title (SDITITL)
- List Libraries (SDILLST)
- List Library Members (SDILMBR)
- Message Inquiry (SDIINQM)
- Message Queue Query (SDIQUMS)
- Pace (SDIPACE)
- Purge Message (SDIPRGM)
- Retrieve Library Member (SDILRTV)
- Retrieve Library Member CDH (SDIRCDH)

- Search Library (SDILSCH)
- Session End (SDISEND)
- Session Inquiry (SDIINQS)
- Session Start (SDISSTA)
- Session Trace Browse (SDITRLB)
- Session Trace Retrieve (SDITRLR)

The Probe (SDIPROB) command can be a response-mode command, depending on whether you request a synchronous or an asynchronous response.



NOTE: The Receive Message command (SDIRCVM) causes Information Exchange to send messages to your system, but is not considered a response-mode command.

Requests that require responses

The following requests to your system from Information Exchange require a response from your address:

- Commit request
- Pace request

When your computer is retrieving messages from Information Exchange, Information Exchange sends Commit or Pace requests to your Information Exchange address to synchronize communication. If you want to checkpoint or synchronize messages that you send, use the Commit or Pace command.

Summary of command sequences

The following explains Information Exchange commands and when to use them:

Session Start	This command must precede all other commands. Information Exchange does not accept any input until you have successfully started a session.
Define Alias	You can usually place this command anywhere within a session; however, it must precede any Send Message or Receive Message command that requires alias definitions.
List Define	You can usually place this command anywhere within a session; however, it must precede any Send Message, Receive Message, or List Verify command referencing the list.
List Verify	You can usually place this command anywhere within a session; however, it must follow any List Define command for the list.

Send Message	<p>This command has several ordering requirements:</p> <ul style="list-style-type: none">• It cannot follow a Commit or Pace command until after the response to the command is received.• A group start must precede a group continuation.• A multiple-segment message must be followed by its continuation segments without interruption by a Pace, Commit, Session End, or Send Message command for different messages.• An <code>ENDATA</code> message ends a prior group.• A message header with an <code>L</code> in <code>MSGGPRIN</code> ends the current message group.• A single message ends a prior group.• A Commit or Pace command does not end a prior group.• A Session End command does not end a prior group.
Audit Browse	<p>With this command, you can retrieve audit information regarding message flow; you can place it anywhere within a session.</p>
Inquiries	<p>You can place inquiry commands anywhere within a session. You can use these commands to make message inquiries.</p>
Queries	<p>You can place query commands anywhere within a session. They are response commands that return information appropriate to the specific query (for example, the Archive Query and Message Queue Query).</p>
Retrievals	<p>These commands cause Information Exchange to move information, such as archived messages, audit information, session traces, and library members into your normal-priority message queue. From there, you can retrieve the information by using the Receive Message command.</p>
Load Test Messages	<p>You can place this command anywhere within a session, but it should precede a Receive Message command to retrieve the test messages.</p>
Cancel	<p>You can place this command anywhere within a session to cancel messages.</p>

Session End	<p>You can place this command anywhere within a session, but if you violate any of the following rules, the session abends:</p> <ul style="list-style-type: none">• Any message group must be completed by sending a single message or an <code>ENDATA</code> message.• A multiple-segment message must have been completed.• If you are using checkpoint-level recovery, no message must have been transmitted except <code>ENDATA</code> messages since the last commit exchange.• Any Receive Message commands must have ended, and the resultant <code>ENDATA</code> must have been sent to your system. This indicates that all Receive Message command processing is complete. If you specify continuous receive in the Receive Message command, you must end it with an <i>end continuous receive</i>. After this is done, continue receiving message segments until your system receives the <code>ENDATA</code>.
Commit	<p>Use this command only in sessions where you have specified checkpoint-level recovery. You can place this command anywhere except between segments of a multiple-segment message.</p>
Pace	<p>You can use this command in any session. You can place it anywhere except between segments of a multiple-segment message.</p>
Probe	<p>This command can provoke an immediate response, depending on whether you request a synchronous or an asynchronous response.</p> <p>Use this command to determine if intended destinations are valid for sending messages with a particular payment method or authorization level, or if the address itself is valid. You can place the Probe command anywhere within a session; however, it must precede any Send Message command to the destinations in question. You cannot use the Probe command for intersecting inquiries.</p>

Means of expansion

With each new release of Information Exchange, there may be new options added to commands. Each command, response, request, or message header contains a command *expansion-level* indicator. As the commands are expanded, new values are assigned to this indicator to show the expansion level of the command.

Information Exchange continues to support previous versions of the command; it is not necessary for you to implement new command versions unless you want to use the new features. Information Exchange lets users with different levels of the commands communicate. It does this by sending requests, responses, and messages in each user's expansion level.

When you update your programming to support the features of a new expansion level, use the same expansion level for all commands and message headers. However, you can handle different expansion levels with Information Exchange. The expansion level for a message header, response, or request that Information Exchange sends to your Information Exchange address is the same as the expansion level from the preceding input to Information Exchange.

Format of commands and responses

This section describes the format of Information Exchange interface commands and responses, including a brief overview of their use. The commands and responses are listed alphabetically for easy reference. However, groups of related commands are kept together. For example, Synchronous Probe response and Asynchronous Probe response are shown immediately following the Probe command as Probe Response (Asynchronous) and Probe Response (Synchronous).

In this section, the term *alphanumeric* generally refers to any keyboard character. However, for practical purposes, use only alphabetic, numeric, space, and common punctuation characters. This is because some terminals or systems might find other characters unacceptable; they might cause the receiving terminal or system to perform erratically. Some systems might interpret special characters differently when they are translated between EBCDIC and ASCII.

Common field name descriptions

Some fields are common to more than one command, response, request, or other item. The following field descriptions apply wherever the field is used.

COMMAND (field 1)

In the case of a command, this field contains a command code that identifies it by its actual name.

In the case of a response, this field contains the value from RETRSPC of the corresponding command.

The format is alphanumeric, left-justified, and padded on the right with blanks.

ACCOUNT (field 2)

This field contains your account ID. The format is alphanumeric, left-justified, and padded on the right with blanks.

ACCNTNO (field 2)

This field contains your account ID. The format is alphanumeric, left-justified, and padded on the right with blanks.

USERID (field 3)

This field contains your user ID. The format is alphanumeric, left-justified, and padded on the right with blanks.

SESSKEY (field 4)

This field must contain the value Information Exchange sends to your system in SESSKEY (field 4) of the Session Start response. The format is alphanumeric, left-justified, and padded on the right with blanks.

Finding descriptions

Commands and responses are grouped into the following categories and listed in alphabetical order within categories:

- “Session commands” on page 87
- “Message transmission commands” on page 118
- “Mailbox commands” on page 142
- “Alias commands” on page 176
- “Distribution List commands” on page 182
- “Library commands” on page 189
- “Audit Commands” on page 231
- “Archive commands” on page 258

Session commands

Session Start command—SDISSTA

This command starts an Information Exchange session. You cannot use other Information Exchange commands until you issue the Session Start command and receive a Session Start response from Information Exchange. The length of the Session Start command is 103 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDISSTA
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	PASSWORD	Password
5	33	8	NEWPSWRD	New password
6	41	8	RETRSPC	Session Start response command name
7	49	8	SYSTYPE	System type
8	57	4	SYSLEVEL	System level
9	61	1	RRLTYPE	Restart/recovery level (<i>X, S, G, C, M</i>)
10	62	5	MAXMSGSZ	Maximum message size
11	67	5	MAXMSGCP	Maximum messages between Commit or Pace requests
12	72	8	COMMCMD	Commit or Pace command code
13	80	8	TESTREST	Test reset indicator
14	88	5	SNDCKPTN	Send checkpoint number
15	93	5	RCVCKPTN	Receive checkpoint number
16	98	5	TIMEZONE	Time zone
17	103	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

PASSWORD (field 4)

This field contains your Information Exchange access password (not to be confused with the network LOGON password). It is used in conjunction with ACCNTNO (field 2) and USERID (field 3) to positively identify you to Information Exchange. The format is alphanumeric, left-justified, and padded on the right with blanks.

NEWPSWRD (field 5)

This field contains either a new password for the next session or blanks if the same password is to be used. A new password becomes effective only after the session ends with a Session End response code of 00000 (normal session end). The format is alphanumeric, left-justified, and padded on the right with blanks. If you are an ESO user, see “Extended security option” on page 11 for information about password restrictions.

RETRSPC (field 6)

This field contains the value the system places in the first 8 characters of the Session Start response. The format is alphanumeric, left-justified, and padded on the right with blanks.

SYSTYPE (field 7)

This field contains a value that identifies your system type. The format is alphanumeric, left-justified, and padded on the right with blanks.

SYSLEVEL (field 8)

This field contains a value that identifies the level or version associated with SYSTYPE (field 7). The format is alphanumeric, left-justified, and padded on the right with blanks.

RRLTYPE (field 9)

This field contains a value that indicates the type of restart and recovery to be used.

This value:	Indicates:
X	Session-level recovery
S	Stream-level recovery
G	Group-level recovery
C	Checkpoint-level recovery
M	Message-level recovery

See “Recovery and restart” on page 14 for more information.

MAXMSGSZ (field 10)

This field contains the maximum message size that your system can receive. MAXMSGSZ is the maximum physical block size that your system can receive, excluding transmission control characters and SNA headers, but including Information Exchange message headers. The maximum message size value should be between 256 and 26,000. If it is over 26,000, 26,000 is used; if it is less than 256, 256 is used. If an expansion level of 3 is used and this field is left blank, the maximum message size used is set to the maximum size that Information Exchange can handle. Information Exchange returns this value to your application in MAXMSGSZ (field 9) of the Session Start response. Information Exchange changes messages that do not fit these limits into multiple-segment messages. See “Transmitting large messages using message segments” on page 25. The format is numeric, right-justified, and padded on the left with zeros.

MAXMSGCP (field 11)

This field contains the maximum number of message segments, excluding ENDDATA messages, your system is prepared to receive between Commit or Pace requests from Information Exchange. The valid range of values for this field depends on the particular installation definitions. Information Exchange does not normally exceed the number specified here, without sending a Commit or Pace request to your system and waiting for a response. If the number of segments required for any single message is greater than this number, Information Exchange does not send the Commit or Pace request to your system for this user ID until the entire message has been sent. The format is numeric, right-justified, and padded on the left with zeros.

COMMCMD (field 12)

This field contains the value that is returned in the first 8 characters of the Commit or Pace requests sent from Information Exchange to the user's system for this user ID during receive processing. The request is a Commit request if the session is using the checkpoint-level or group-level recovery option (*C* or *G* in field 9); otherwise, it is a Pace request. The format is alphanumeric, left-justified, and padded on the right with blanks.

TESTREST (field 13)

RESETSES and RESETMSG values are used only if you are still logically in session with Information Exchange (no successful session end has been performed).

This value:	Indicates:
blank	A normal session start or restart.
RESETSES	Reset the session. The session is a cold session start (as if there were no preceding session), as opposed to a normal restart (some fields of the prior session are retained in the restarted session). Any messages on the message queue for your user ID are reset so they can be received.
RESETMSG	Reset the session and erase the messages. The session is restarted as a cold session start, as opposed to a normal restart. Any messages on the message queue for your user ID are deleted.



NOTE: The ability to use this field is determined by a flag in the user-profile record. This flag can be changed through Information Exchange Administration Services.

SNDCKPTN (field 14)

This field contains the last completed checkpoint number for sending messages. If checkpoint-level or group-level recovery is not used, Information Exchange ignores this field. If a new checkpoint-level or group-level recovery session is being started, as opposed to being restarted after a failure, specify *00000* in this field, or leave it blank. The format is numeric, right-justified, and padded on the left with zeros.

RCVCKPTN (field 15)

This field contains the last completed checkpoint number for receiving messages. If checkpoint-level or group-level recovery is not used, Information Exchange ignores this field. If a new checkpoint-level or group-level recovery session is being started (as opposed to being restarted after a failure), specify *00000* in this field, or leave it blank. The format is numeric, right-justified, and padded on the left with zeros.

TIMEZONE (field 16)

This field contains the time zone to which you would like the dates and times offset. Unless otherwise specified, this means that any date or time returned by Information Exchange during this session will be in the time zone specified in this field. Specify the time zone in the format *Ehhmm*, where *hh* indicates the number of hours and *mm* the number of minutes east of the Greenwich meridian, or *Whhmm*, indicating the hours and minutes west of the Greenwich meridian.

Alternatively, you might prefer to use one of the following time zone abbreviations. The format is left-justified and padded on the right with blanks.

Time zone	GMT offset	Description
EAD	E1000	Eastern Australia daylight
JST	E0900	Japan standard time
WED	E0200	Western Europe daylight
EMT	E0200	Eastern Mediterranean time
BST	E0100	British summer time
WES	E0100	Western Europe standard
GMT	E0000	Greenwich mean time
AST	W0400	Atlantic standard time
EDT	W0400	Eastern daylight time
EST	W0500	Eastern standard time
CDT	W0500	Central daylight time
CST	W0600	Central standard time
MDT	W0600	Mountain daylight time
MST	W0700	Mountain standard time
PDT	W0700	Pacific daylight time
PST	W0800	Pacific standard time
YDT	W0800	Alaska daylight time
YST	W0900	Alaska standard time
AHD	W1000	Hawaii standard time
AHS	W1000	Hawaii standard time

EXPAND (field 17)

This field contains a command expansion-level indicator.

This value: Indicates:

- blank or 1 The Session Start response data is to return fields 1 through 8.
- 2 Reserved.
- 3 The Session Start response data is to return fields 1 through 11.

Session Start response

This is Information Exchange's response to the Session Start command; its length is 57 bytes for expansion level 3 or 48 bytes for expansion level 1.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDISSTA
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	5	RESPRSN	Session Start response code
6	38	5	LSTSNDCP	Last send checkpoint number
7	43	5	LSTRCVCP	Last receive checkpoint number
8	48	1	EXPAND	Expansion indicator
9	49	5	MAXMSGSZ	Maximum message size
10	54	2	VERSION	Information Exchange version number
11	56	2	RELEASE	Information Exchange release number

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

RESPRSN (field 5)

This field contains the Session Start response code, which is 00000 if the session started successfully. See "Session Start response codes" on page 93 for explanations of other values.

LSTSNDCP (field 6)

This field contains the last completed checkpoint number that Information Exchange assigned while you were sending messages to Information Exchange. The format is numeric, right-justified, and padded on the left with zeros.

For checkpoint-level or group-level recovery, this field contains the last checkpoint number assigned by Information Exchange in a Commit response. If no commit point for input operations has been reached, the value is zero. If the value does not match what your system expected (as sent in the Session Start command, field 14), your system retransmits messages sent to Information Exchange since the indicated checkpoint number in this message.

If stream-level recovery is used, a value of 1 in this field indicates that Information Exchange did not commit a stream of messages sent in the prior session. A value of zero indicates that no stream of messages sent to Information Exchange was successfully processed.

For message- and session-level recovery, this field is set to zero. For more information, see "Recovery and restart" on page 14.

LSTRCVCP (field 7)

This field contains the last completed checkpoint number that your system sent to Information Exchange for the indicated user ID in a Commit response while receiving messages. The format is numeric, right-justified, and padded on the left with zeros.

For checkpoint-level or group-level recovery, this field contains the last checkpoint number your system assigned for committed messages sent from Information Exchange. If no commit point for output operations has been reached, this is zero. If the value does not match what your system expected, as sent in the Session Start command, your system discards messages it received since this checkpoint number.



NOTE: If a session is restarted using checkpoint-level or group-level recovery and messages are being sent to your system, they are resent following the Session Start response.

If stream-level recovery is used, a value of 1 in this field indicates that Information Exchange has not reached the end of the output phase of the prior session, and that messages that were sent to your system are sent again.

EXPAND (field 8)

This field contains the value from EXPAND (field 17) of the Session Start command.

MAXMSGSZ (field 9)

This field contains the maximum message size that Information Exchange uses. This value is the maximum physical block size that your system can transmit, excluding transmission control characters and SNA headers.

VERSION (field 10)

This field contains the Information Exchange version number.

RELEASE (field 11)

This field contains the Information Exchange release number.



NOTE: The system sends only MAXMSGSZ (field 9), VERSION (field 10), and RELEASE (field 11) if the value of EXPAND (field 8) in the Session Start command is 3 or higher.

Session Start response codes

The following lists and describes the values of field RESPRSN in the Session Start response. If the value in RESPRSN is less than *00064*, the session started. If the value of RESPRSN is greater than *00063*, the session did not start.

Interpret the value Information Exchange assigns to RESPRSN as a binary number, in which each power of two represents a particular condition. Information Exchange assigns each condition a value of *1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192, or 16384*. When more than one condition exists, the system represents it as the sum of the values of the individual conditions. For example, *21* represents values $16 + 4 + 1$, showing that the conditions indicated by values *16, 4, and 1* exist.

This value:	Indicates:
00000	Normal condition; no error exists. Information Exchange started the session exactly as you requested.
00001	In a session using checkpoint-level recovery, the checkpoint numbers for the send and receive side of the session do not match Information Exchange's values. LSTSNDPCP and LSTRCVCP indicate the numbers at which Information Exchange is restarting.
00002	The system is restarting the session. There was a prior session, and it is being restarted.
00004	The time zone field is invalid. Information Exchange assumes the value of the time zone, and stores a message you can retrieve specifying the assumed time zone.
00008	The value of MAXMSGCP is invalid. The valid range of values for this field depends on the particular installation definition.
00016	The value of MAXMSGSZ is invalid. The valid range of values for this field depends on the particular installation definition.
00032	The value of TESTREST is invalid. Information Exchange ignores your request to reset the session and the queued messages.

This value:	Indicates:
00064	<p>If you are an ESO user, you specified a new password, NEWPSWRD (field 5), on the Session Start command record that did not meet the following ESO password rules:</p> <ul style="list-style-type: none">• Must not contain the user ID as any part• Must be at least six characters in length• Must contain at least three different characters• Must contain a nonnumeric first and last character• Must contain at least one non-alphabetic character• Must contain at least one alphabetic character• Must contain only the valid characters A-Z, 0-9, and special characters # @ and \$• Must not contain more than two identical consecutive characters• Must be different from the current or five previous passwords• Must not contain more than three identical consecutive characters from the previous password <p>Change your password to conform to these rules.</p>
00128	<p>The value of RCVCKPTN is invalid. The session does not start. RCVCKPTN must be blanks (if not restarting receive from a checkpoint) or numeric.</p>
00256	<p>The value of SNDCKPTN is invalid. The session does not start. SNDCKPTN must be blanks (if not restarting send from a checkpoint) or numeric.</p>
00512	<p>The restart recovery level differs from the original recovery level. The session does not start. The recovery level must be the same in any session restart as in the original session start.</p>
01024	<p>The value of RRLTYPE is invalid. The session does not start. RRLTYPE must be <i>C</i>, <i>M</i>, <i>S</i>, <i>X</i>, or <i>G</i>.</p>
02048	<p>The password is incorrect. The session does not start. The password for this Information Exchange address is not the same as specified in your Session Start command.</p>
04096	<p>The user ID is invalid. The session does not start. The address is not known to Information Exchange.</p>
08192	<p>If you are an ESO user, you have sent three successive Session Starts using an incorrect password. Your Information Exchange user ID has been revoked. Contact your service administrator to request that your password be reset using Information Exchange Administration Services. Resetting the password resumes the user ID.</p>

This value:	Indicates:
16384	If you are an ESO user, you did not specify a new password on the Session Start command. If the Information Exchange password for an ESO user is the same as that user's normal user ID, a new password must be specified. Modify NEWPSWRD (field 5 in the Session Start command) to contain a password that conforms to the rules described for the Session Start response code value of <i>64</i> .

Session End command—SDISEND

This command ends a session normally. Its length is 46 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDISEND
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	RETRSPC	Session end response command
6	41	5	SENDCODE	Session end code
7	46	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

RETRSPC (field 5)

This field contains the value that is returned in the first 8 characters of the Session End response. The format is alphanumeric, left-justified, and padded on the right with blanks.

SENDCODE (field 6)

If this field is not zeros, the session ends abnormally, permitting a subsequent session restart from the most recently established recovery point. To end the session normally, specify zeros in this field. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 7)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

Session End response

This is Information Exchange's response to the Session End command; its length is 38 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDISEND
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	5	SENDCODE	Session end condition code
6	38	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

SENDCODE (field 5)

This field contains a value that indicates the results of processing the Session End command; *00000* indicates the session ended normally. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 6)

This field contains the value from EXPAND (field 7) of the Session End command.

Session End response codes

The following lists and describes the values of field SENDCODE in the Session End response. If the value in SENDCODE is zero, the session ended normally. Otherwise, the session did not end normally, and you should restart it.



NOTE: If you specified a new password in the preceding Session Start command, it becomes active only if the session ends normally. This occurs when you receive a value of *00000* in the Session End response.

Interpret the value Information Exchange assigns to SENDCODE as a binary number, in which each power of two represents a particular condition. Information Exchange assigns each condition a value *1, 2, 4, 8, 16, or 32*. When more than one condition exists, the system represents it as the sum of the values of the individual conditions. For example, *21* represents values *16 + 4 + 1*, showing that the conditions indicated by values *16, 4, and 1* all exist.

This value:	Indicates:
00000	Normal condition; no error exists. Information Exchange ended the session correctly.
00001	You have requested an abnormal end by placing a value other than zeros in SENDCODE (field 6) of the Session End command.
00002	The session was waiting for a Commit or Pace response. Information Exchange is sending messages to your system, and had sent a Commit or Pace request, but Information Exchange received the Session End command before the Commit or Pace response. Make certain your program generates correct command sequences.
00004	One or more Receive Message commands had not ended when Information Exchange received the Session End command. Do not send the Session End command until you receive an ENDDATA for each Receive Message command you issue to Information Exchange.
00008	An incomplete group of messages or uncommitted messages was sent to Information Exchange. An ENDDATA message should be sent before the Session End command to show that the last message group sent to Information Exchange is completed. If you are using checkpoint-level or group-level recovery and there are uncommitted messages, a Commit command should be issued before the Session End command.
00016	Wrong access key. The Session End command did not contain the access key assigned to the session when it was started.
00032	No session was in progress.

Session Inquiry command—SDIINQS

This command obtains status information from Information Exchange about a specific session. You can request a long- or short-format response. The length of this command is 42 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIINQS
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	RETRSPC	Inquiry response command
6	41	1	INQTYPE	Inquiry type (<i>S, L</i>)
7	42	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

RETRSPC (field 5)

This field contains a value that is returned in the first 8 characters of the Session Inquiry response (both long and short format). The format is alphanumeric, left-justified, and padded on the right with blanks.

INQTYPE (field 6)

This field contains a value that determines the extent of the information returned in response to this command.

This value:	Indicates:
S	A short-form response is requested.
L	A long-form response is requested.

EXPAND (field 7)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

Session Inquiry response

This is Information Exchange's response to the Session Inquiry command, if the long-format response was requested; its length is 190 bytes; if the short-format response was requested, its length is 85 bytes.

Long format

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIINQS
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	INQTYPE	Inquiry type
6	34	6	INQDATE	Inquiry date
7	40	6	INQTIME	Inquiry time
8	46	6	SESSDATE	Session start date
9	52	6	SESSTIME	Session start time
10	58	6	SESSLLEN	Session duration
11	64	6	SNDTOT	SDISNDM commands sent
12	70	2	ACTVRCV	SDIRCVM commands active
13	72	6	RCVTOT	Total messages received
14	78	6	IDLETIME	Idle state duration
15	84	1	SYSWAIT	SDIPACE/SDICMIT response indicator
16	85	6	SNDPROC	SDISNDM commands processed
17	91	6	SNDAOK	SDISNDM commands accepted
18	97	6	SNDREJ	SDISNDM commands rejected
19	103	6	RCVDATA	Data messages received
20	109	6	RCVSERV	Services messages received
21	115	6	RCVOTHR	Other messages received
22	121	5	LSTSNDCP	Last send checkpoint number
23	126	5	LSTRCVCP	Last receive checkpoint number
24	131	29	LSTMSGs	Last data message sent
25	160	30	LSTMSGr	Last data message received
26	190	1	EXPAND	Expansion indicator

Short format

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIINQS
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	INQTYPE	Inquiry type
6	34	6	INQDATE	Inquiry date
7	40	6	INQTIME	Inquiry time
8	46	6	SESSDATE	Session start date
9	52	6	SESSTIME	Session start time
10	58	6	SESSLN	Session duration
11	64	6	SNDTOT	Total SDISNDM commands sent
12	70	2	ACTVRCV	Total SDIRCVM commands active
13	72	6	RCVTOT	Total messages received
14	78	6	IDLETIME	Idle state duration
15	84	1	SYSWAIT	System waiting indicator
16	85	1	EXPAND	Expansion indicator

COMMAND (field 1)**ACCNTNO (field 2)****USERID (field 3)****SESSKEY (field 4)****INQTYPE (field 5)**

This field contains a value that indicates the format of this response.

This value:	Indicates:
L	Long format
S	Short format

INQDATE (field 6)

This field contains the date the inquiry was processed, offset to the time zone specified in the Session Start command. The format is numeric (*YYMMDD*), right-justified, and padded on the left with zeros.

INQTIME (field 7)

This field contains the time the inquiry was processed, offset to the time zone specified in the Session Start command. The format is numeric (*HHMMSS*), right-justified, and padded on the left with zeros.

SESSDATE (field 8)

This field contains the date the session was started, offset to the time zone specified in the Session Start command. The format is numeric (*YYMMDD*), right-justified, and padded on the left with zeros.

SESSTIME (field 9)

This field contains the time the session was started, offset to the time zone specified in the Session Start command. The format is numeric (*HHMMSS*), right-justified, and padded on the left with zeros.

SESSLEN (field 10)

This field contains the duration of this session. The format is numeric (*HHMMSS*), right-justified, and padded on the left with zeros.

SNDTOT (field 11)

This field contains the total number of the Send Message command message segments that Information Exchange receives for this session. The format is numeric, right-justified, and padded on the left with zeros.

ACTVRCV (field 12)

This field contains the total number of currently active Receive Message commands for this session. The format is numeric, right-justified, and padded on the left with zeros.

RCVTOT (field 13)

This field contains the total number of message segments sent to the user during the current session as a result of Receive Message commands. The format is numeric, right-justified, and padded on the left with zeros.

IDLETIME (field 14)

This field contains the time lapsed since the last message transfer activity (receipt of a Send Message command, output of a message to satisfy a Receive Message command, or a Pace or Commit command) in the current session. The format is numeric (*HHMMSS*), right-justified, and padded on the left with zeros.

SYSWAIT (field 15)

This field contains a value that indicates whether the session is waiting for a Commit response, a Pace response, or neither.

This value:	Indicates:
P	Information Exchange is waiting for a Pace response from your user ID.
C	Information Exchange is waiting for a Commit response from your user ID.
blank	Information Exchange is not waiting for either the Pace or Commit response.

SNDPROC (field 16)

This field contains the total number of Send Message commands (complete messages, not message segments) processed by Information Exchange for this session. Information Exchange processes Send Message commands only after they have reached a recovery point. The format is numeric, right-justified, and padded on the left with zeros.

SNDAOK (field 17)

This field contains the total number of Send Message commands (complete messages, not message segments) processed and accepted for delivery service by Information Exchange for this session. The format is numeric, right-justified, and padded on the left with zeros.

SNDREJ (field 18)

This field contains the total number of Send Message commands (complete messages, not message segments) processed, but rejected, by Information Exchange for this session. The format is numeric, right-justified, and padded on the left with zeros.

RCVDATA (field 19)

This field contains the total number of user data messages (complete messages, not message segments) sent to the user as a result of Receive Message commands during the current session. The format is numeric, right-justified, and padded on the left with zeros.

RCVSERV (field 20)

This field contains the total number of Information Exchange system messages (complete messages, not message segments) sent to the user as a result of Receive Message commands during the current session. The format is numeric, right-justified, and padded on the left with zeros.

RCVOTHR (field 21)

This field contains the total number of non-system messages queued to the user as a result of Information Exchange commands, such as the Session Inquiry command, during the current session. The format is numeric, right-justified, and padded on the left with zeros.

LSTSNDCP (field 22)

This field contains the last send checkpoint number established with Information Exchange during the current session. The format is numeric, right-justified, and padded on the left with zeros.

LSTRCVCP (field 23)

This field contains the last receive checkpoint number established with Information Exchange during the current session. The format is numeric, right-justified, and padded on the left with zeros.

LSTMSGs (field 24)

This field contains the values of DESTACCT (field 8), DESTUID (field 9), MSGNAME (field 11), and MSGSEQN (field 12) from the last Send Message command received by Information Exchange during the current session. The format is alphanumeric, left-justified, and padded on the right with blanks.

LSTMSGR (field 25)

This field contains the value of the network message identification (an internal control field for Information Exchange) and MSGSEQO (field 18) in the last message sent in response to a Receive Message command during the current session. The format is alphanumeric, left-justified, and padded on the right with blanks.

EXPAND (field 26--long format, or field 16--short format)

This field contains the value from EXPAND (field 7) of the Session Inquiry command.

Session Trace Browse command—SDITRLB

You can use this command to browse the contents of your session trace data. Information Exchange's response to this command is the Session Trace Browse response. This response contains the session trace records that provide information on the status of your session. The length of this command is 98 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDITRLB
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	1	RECTYPES	Type of records to retrieve (<i>A, T, S, R</i>)
8	43	8	ALTACCT	Alternate account ID
9	51	8	ALTUSRID	Alternate user ID
10	59	1	TIMEZONE	Time zone
11	60	7	DATEFROM	From date (CYYMMDD)
12	67	7	DATETO	To date (CYYMMDD)
13	74	25	LOCATOR	Command chaining value

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to 1.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Session Browse Trace response.

RECTYPES (field 7)

This field specifies the type of session trace records Information Exchange returns. The format is a coded value.

This value: Indicates:

T Both session trace and session trace receive records. The following are valid subsets of *T*:

This value:

Indicates:

S

Session trace records.

R

Session trace receive records.

A

All record types. This is the default.

ALTACCT (field 8)

This field allows you to request session records for an account/user ID other than your own. If left blank, session records for the account specified in ACCNTNO (field 2) are returned. If a value is specified in this field, ALTUSRID (field 9) must also be specified. The requester must have authority to view session records for the alternate account/user ID.

ALTUSRID (field 9)

This field, together with ALTACCT (field 8), allows you to request session records for an account/user ID other than your own or for a cluster parent user ID. You must be authorized to request session records other than your own. Information Exchange ignores input from nonauthorized users. If you leave this field blank or you are not authorized to view session records for the alternate account/user ID, Information Exchange extracts only your own session records. If a value is specified in this field but the ALTACCT (field 8) is left blank, Information Exchange extracts audit records for the account specified in the ACCNTNO (field 2) and the user ID specified in this field. To extract session records for all users in the specified account, enter a question mark (?) in this field.

TIMEZONE (field 10)

This field indicates the time zone for the dates in DATEFROM (field 10) and DATETO (field 11). Enter either the value L (for the time zone value specified in the Session Start command) or the value G (for GMT). The default is L.

DATEFROM (field 11)

Information Exchange extracts only session trace records representing events on or after the date you enter in this field. Enter the date in the format *CYYMMDD*, where *C* is the century (0=1900, 1=2000), *YY* is the year, *MM* is the month, and *DD* is the day of the month. For example, 0890502 represents May 2, 1989. The default is 0000102 (January 2, 1900).

DATETO (field 12)

Information Exchange extracts only session trace records representing events on or before the date you enter in this field. Enter the date in the format *CYYMMDD*, where *C* is the century (0=1900, 1=2000), *YY* is the year, *MM* is the month, and *DD* is the day of the month. For example, 1010502 represents May 2, 2001. The default is 1420916 (September 16, 2042).

LOCATOR (field 13)

This field contains a command chaining value that receives information for additional session trace records. Because of the maximum message size specification you indicated in MAXMSGSZ (field 10) of the Session Start command, Information Exchange may not have been able to send you all the available session trace records in one block. Use this field to indicate that you want to browse the records from the last breakpoint. Submit the initial Session Trace Browse command with blanks in this field. Submit subsequent Session Trace Browse commands with the CHAINRSP (field 6) value from the previous Session Trace Browse response in this field, until you receive a response with a series of 9s in CHAINRSP.

Session Trace Browse response

This is Information Exchange's response to the Session Trace Browse command. This response contains the session trace records that provide information on the status of your session. The length of the Session Trace Browse response is 63 bytes, plus the length of the items returned.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDITRLB
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	25	CHAINRSP	Command chaining value
7	59	2	ERRCODE	Error code
8	61	3	ITEMS	Number of trace items to follow
9	64	n	item one	Beginning of first session trace item

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator from EXPAND (field 5) of the Session Trace Browse command.

CHAINRSP (field 6)

This field contains a command chaining value to be copied into LOCATOR (field 13) of the subsequent Session Trace Browse command. If this value is a series of 9s, there are no more records that meet your retrieval specifications.

ERRCODE (field 7)

A value of *00* in this field indicates Information Exchange did not find any errors while processing the command. A value of *04* indicates there was a recoverable error. A value of *08* indicates there was a nonrecoverable error present. If there is any value other than *00*, issue a Receive Message command to retrieve any error messages.

ITEMS (field 8)

This field contains the number of trace entries to follow.

session trace items (field 9)

This field contains the session trace records. "Session Trace Receive Record—character format" on page 111 and "Session Trace Record—character format" on page 115 for a description of the session trace record formats.

Session Trace Retrieve command—SDITRLR

This command downloads the contents of your session trace file into your Information Exchange mailbox in the form of a message group. When this is complete, you can receive the messages with the Receive Message command. If the system does not find session trace records that meet your criteria, it places a null (zero length) message in your mailbox. The length of this command is 76 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDITRLR
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	1	RECTYPES	Type of records to retrieve (<i>A, T, S, R</i>)
8	43	8	ALTACCT	Alternate account ID
9	51	8	ALTUSRID	Alternate user ID
10	59	1	TIMEZONE	Time zone
11	60	7	DATEFROM	From date (CYYMMDD)
12	67	7	DATETO	To date (CYYMMDD)
13	74	3	MAXMSGSZ	Maximum message length

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to 1.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Session Trace Retrieve response.

RECTYPES (field 7)

This field specifies the type of session trace records Information Exchange returns. The format is a coded value.

This value: Indicates:

T Both session trace and session trace receive records. The following are valid subsets of *T*:

This value: Indicates:

S Session trace records.

R Session trace receive records.

A All record types. This is the default.

ALTACCT (field 8)

This field allows you to request session records for an account/user ID other than your own. If left blank, session records for the account specified in **ACCOUNT (field 2)** are returned. If a value is specified in this field, **ALTUSRID (field 8)** must also be specified. The requester must have authority to view session records for the alternate account/user ID.

ALTUSRID (field 9)

This field, together with **ALTACCT (field 8)**, allows you to request session records for an account/user ID other than your own or for a cluster parent user ID. You must be authorized to request session records other than your own. Information Exchange ignores input from nonauthorized users. If you leave this field blank or you are not authorized to view session records for the alternate account/user ID, Information Exchange extracts only your own session records. If a value is specified in this field but the **ALTACCT (field 8)** is left blank, Information Exchange extracts session records for the account specified in the **ACCOUNT (field 2)** and the user ID specified in this field. To extract session records for all users in the specified account, enter a question mark (?) in this field.

TIMEZONE (field 10)

This field contains a value that indicates the time zone for the dates in **DATEFROM (field 10)** and **DATETO (field 11)**. Enter either the value *L* (for the time zone value specified in the Session Start command) or the value *G* (for GMT). The default is **L**.

DATEFROM (field 11)

Information Exchange extracts only session trace records representing events on or after the date you enter in this field. Enter the date in the format *CYYMMDD*, where *C* is the century (0=1900, 1=2000), *YY* is the year, *MM* is the month, and *DD* is the day of the month. For example, 0890502 represents May 2, 1989. The default is 0000102 (January 2, 1900).

DATETO (field 12)

Information Exchange extracts only session records representing events on or before the date you enter in this field. Enter the date in the format *CYYMMDD*, where *C* is the century (0=1900, 1=2000), *YY* is the year, *MM* is the month, and *DD* is the day of the month. For example, 1010502 represents May 2, 2001. The default is 1420916 (September 16, 2042).

MAXMSGSZ (field 13)

This field indicates the size of the largest message to be created in the message group. The value must be numeric. This field enables checkpoint-level recovery to take commit points while receiving the message group that Information Exchange builds. If MAXMSGSZ contains zero or any nonnumeric characters, Information Exchange creates a single message in the message group. The value in MAXMSGSZ is multiplied by 1000 to obtain the message size.

Output Format

The Session Trace Retrieve command builds a message group in your Information Exchange mailbox. The following fields are in the message group's identification fields:

This field:	Contains:
DESTACCT	*SYSTEM*
DESTUID	*EVENTS*
DESTTYPE	D
MSGNAME	blanks
MSGSEQN	blanks
MSGNCLS	blank
MSGUCLS	cccccccc
MSGCLASS	blanks
MSGRCPTS	blank
MSGCHRG	5
SYSTYPE	IBMIE
SYSLEVEL	Information Exchange version and release level
DTBLTYP	blank
DTBLID	blanks



NOTE: The value in MSGUCLS identifies the type of information retrieved. The following are valid sub-values of MSGUCLS.

This value:	Retrieves:
#STRACE	All session trace information, including both session and receive trace data. This value is used if you placed a <i>T</i> or an <i>A</i> in RECTYPES (field 7).
#SESSION	Only session trace data. This value is used if you placed an <i>S</i> in RECTYPES (field 7).
#RECEIVE	Only receive trace data. This value is used if you placed an <i>R</i> in RECTYPES (field 7).

The content of the logical record in the message group is illustrated in “Session Trace Receive Record—character format” on page 111 and “Session Trace Record—character format” on page 115.

Session Trace Retrieve response

This is Information Exchange's response to the Session Trace Retrieve command; its length is 35 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value for RETRSPC of SDITRLR
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Response code

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator from EXPAND (field 5) of the Session Trace Retrieve command.

RETCODE (field 6)

If this field contains a value other than 00, the system encountered one or more errors during processing. This is the maximum error message severity code. To find out the cause of an error, receive your Information Exchange error messages from your mailbox and look at their content. If a code less than 08 displays, the system places session trace records in your Information Exchange mailbox.

Session Trace Receive Record—character format

The following record structure represents the format of the Session Trace Receive records as they exist within the Information Exchange message or the Session Trace Browse response. All the fields are in character format. The length of the Session Trace Receive record is 123 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	16	RUSERID	Account ID/User ID
2	17	6	RDATE	Receive start date
3	23	6	RTIME	Receive start time
4	29	1	RTYPE	Type of trace record
5	30	15	RTPHRASE	Type of record (phrase)
6	45	8	RSESSKEY	Session access key
7	53	8	RREFNAME	SDIRCVN reference ID
8	61	1	RALIAST	Source alias table type
9	62	3	RALIASTB	Source alias table name
10	65	8	RSRCACCT	Source account
11	73	8	RSRCUID	Source user ID
12	81	1	RSRCTYP	Type of source ID
13	82	8	RUCLASS	Message user class
14	90	1	RCMDTYP	Type of receive
15	91	1	REXPAND	Expansion indicator
16	92	8	RARCREF	Archive reference
17	100	20	RMSGKEY	Specific receive message key
18	120	3	reserved	Reserved
19	123	1	REDI	EDI indicator

RUSERID (field 1)

This field contains the Information Exchange address of the user who requested the receive. The first 8 characters represent the user's Information Exchange account; the last 8 characters represent the user's ID.

RDATE (field 2)

This field contains the date the receive started, offset to the time zone specified in the Session Trace Browse command or Session Trace Retrieve command. The format is *YYMMDD*.

RTIME (field 3)

This field contains the time the receive started, offset to the time zone specified in the Session Trace Browse command or Session Trace Retrieve command. The format is *HHMMSS*.

RTYPE (field 4)

This field contains the number 3, which indicates a Session Trace Receive record.

RTPHRASE (field 5)

This field contains a short phrase indication of one of the above type values.

RSESSKEY (field 6)

This field contains the value Information Exchange sent to the receiver system in SESKEY (field 4) of the Session Start response. The format is alphanumeric, left-justified, and padded on the right with blanks.

RREFNAME (field 7)

This field identifies this Receive Message command and distinguishes this command from others that were simultaneously active.

RALIAST (field 8)

See “Understanding source address specification” on page 113.

RALIASTB (field 9)

See “Understanding source address specification” on page 113.

RSRCACCT (field 10)

See “Understanding source address specification” on page 113.

RSRCUID (field 11)

See “Understanding source address specification” on page 113.

RSRCTYP (field 12)

This describes the type of source address specification used.

This value: Indicates:

blank	The command received any messages queued to the user’s ID, regardless of the originating user ID.
D	The command received any messages queued to the user’s ID from a specific originating user ID. A local true ID, an intersystem ID, or an alias reference was used.
L	The name of a list of user IDs from which you want to receive messages.
A	The command received all archived messages requeued to the user’s normal-priority message queue.

See “Understanding source address specification” on page 113.

RUCLASS (field 13)

This field contains the message user class the Receive Message command specified.

RCMDTYP (field 14)

This field contains the type of Receive Message command issued.

This value: Indicates:

blank	A batch-receive operation.
C	A continuous-receive operation.
E	A request to end a prior continuous-receive operation.
G	A single group receive operation.
W	A single group receive operation with the option to wait.

REXPAND (field 15)

This field contains the value from EXPAND on the Receive Message command.

RARCREF (field 16)

This field contains the value assigned to the messages that are stored in the short-term archive.

RMSGKEY (field 17)

This field contains the message group key specified in the Receive Message command.

RESERVED (field 18)

This field is reserved.

REDI (field 19)

This field contains the EDI indicator the Receive Message command specified

Understanding source address specification

The value in RSRCTYP (field 12) indicates the address specification used. A value of *D* in this field indicates a local true ID, an intersystem ID, or an alias reference. A value of *L* indicates a distribution list.

A local true ID is specified as follows:

Field:	Contains:
RSRCTYP (field 12)	<i>D</i>
RALIAST (field 8)	blank
RALIASTB (field 9)	blank
RSRCACCT (field 10)	Source account ID
RSRCUID (field 11)	Source user ID

An intersystem ID is specified as follows:

Field:	Contains:
RSRCTYP (field 12)	<i>D</i>
RALIAST (field 8)	<i>I</i>
RALIASTB (field 9)	Intersystem system ID
RSRCACCT (field 10)	Source account ID
RSRCUID (field 11)	Source user ID

An alias reference is specified as follows:

Field:	Contains:
RSRCTYP (field 12)	<i>D</i>
RALIAST (field 8)	Table type (<i>G, O, or P</i>)
RALIASTB (field 9)	Table name
RSRCACCT (field 10)	First 8 characters of the alias name
RSRCUID (field 11)	Last 8 characters of the alias name

A distribution list is specified as follows:

Field:	Contains:
RSRCTYP (field 12)	<i>L</i>
RALIAST (field 8)	blank
RALIASTB (field 9)	blank
RSRCACCT (field 10)	Distribution list name
RSRCUID (field 11)	blank

Session Trace Record—character format

The following record structure is the format of the Session Trace records as they exist within the Information Exchange message or the Session Trace Browse response. All the fields are in character format. The length of the Session Trace record is 148 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	16	SUSERID	Account ID/User ID
2	17	6	SSDATE	Session start date
3	23	6	SSTIME	Session start time
4	29	1	STYPE	Type of trace record
5	30	15	STPHRASE	Type of record (phrase)
6	45	8	SSESSKEY	Session access key
7	53	8	SSYSTYPE	System type
8	61	4	SSYSLVL	System level
9	65	1	SRECLVL	Recovery level
10	66	5	SSRCODE	Session Start response code
11	71	6	SENDDATE	Session end date
12	77	6	SENDTIME	Session end time
13	83	5	SENDCODE	Session end response code
14	88	5	SMSGSIZE	Transmission size
15	93	6	SSNDGRPS	Number of groups sent
16	99	6	SRCVGRPS	Number of groups received
17	105	1	SEXPAND	Expansion level
18	106	5	SSCPEX	Send checkpoint expected
19	111	5	SSCPCV	Send checkpoint received
20	116	5	SRCPEX	Receive checkpoint expected
21	121	5	SRCPCV	Receive checkpoint received
22	126	15	SENDREAS	Reason for session end
23	141	8	SAPPLID	Application ID for session

SUSERID (field 1)

This field contains the Information Exchange address of the user who started the session. The first 8 characters represent the user's Information Exchange account; the last 8 characters represent the user's ID.

SSDATE (field 2)

This is the date on which the session started, offset to the time zone specified in the Session Trace Browse command or Session Trace Retrieve command. The format is *YYMMDD*.

SSTIME (field 3)

This is the time when the session started, offset to the time zone specified in the Session Trace Browse command or Session Trace Retrieve command. The format is *HHMMSS*.

STYPE (field 4)

This field contains the value 0, which indicates a Session Trace record.

STPHRASE (field 5)

This is a short phrase indication of the STYPE (field 4) value.

SSESSKEY (field 6)

This field contains the value Information Exchange sent to the user's system in `SESSKEY` (field 4) of the Session Start response. The format is alphanumeric, left-justified, and padded on the right with blanks.

SSYSTYPE (field 7)

This field contains the system type of the user who started the session with Information Exchange.

SSYSLVL (field 8)

This field contains the system level of the user who started the session with Information Exchange.

SRECLVL (field 9)

This value identifies the restart/recovery level of this session.

This value:	Indicates:
C	Checkpoint-level recovery.
G	Group-level recovery.
M	Message-level recovery.
X	Session-level recovery.
S	Stream-level recovery.

SSRCODE (field 10)

This field contains the Session Start Response code `00000` if the session started successfully.

SENDDATE (field 11)

This is the date when the session ended, offset to the time zone specified in the Session Trace Browse command or Session Trace Retrieve command. The format is `YYMMDD`. It contains blanks if the session did not end with a Session End command.

SENDTIME (field 12)

This is the time when the session ended, offset to the time zone specified in the Session Trace Browse command or Session Trace Retrieve command. The format is `HHMMSS`. It contains blanks if the session did not end with a Session End command.

SENDCODE (field 13)

This field contains the Session End Response code `00000` if the session ended normally. If the session did not end with a Session End command, it contains blanks.

SMSGSIZE (field 14)

This field contains the value of the maximum message size the user's system can receive, as specified on the Session Start command.

SSNDGRPS (field 15)

This field contains the number of message groups sent during this session.

SRCVGRPS (field 16)

This field contains the number of message groups received during this session.

SEXPAND (field 17)

This field contains the value from EXPAND on the Session Start command.

SSCPEX (field 18)

This field contains the checkpoint number expected by the user's system for sending messages.

SSCPCV (field 19)

This field contains the last completed checkpoint number assigned by Information Exchange for sending messages.

SRCPEX (field 20)

This field contains the checkpoint number expected by the user's system for receiving messages.

SRCPCV (field 21)

This field contains the last completed checkpoint number assigned by Information Exchange for receiving messages.

SENDREAS (field 22)

This field contains the session end reason descriptive phrase.

SAPPLID (field 23)

This field contains the application ID (APPLID) that started the session.

Message transmission commands

Cancel command—SDICNCL

This command uniquely identifies messages you want to cancel and causes Information Exchange to delete messages you specify from the destination's queue.

If the expansion level is equal to 1, then the length of this command is 102 bytes and contains fields 1-20. If the expansion level is equal to 2, then this command is 104 bytes and contains fields 1 - 22.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDICNCL
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	DESTACCT	Destination account or distribution list name or first 8 characters of destination alias name
6	41	8	DESTUID	Destination user ID or last 8 characters of destination alias name
7	49	1	DESTTYPE	Destination type (<i>D, L</i>)
8	50	1	MSGCLASS	Message delivery class (<i>blank, P</i>)
9	51	8	MSGNAME	Message name
10	59	5	MSGSEQN	Message sequence number
11	64	8	MSGUCLS	User message classification
12	72	6	SUBDSTA	Submit start date (YYMMDD)
13	78	6	SUBTSTA	Submit start time (HHMMSS)
14	84	6	SUBDEND	Submit end date (YYMMDD)
15	90	6	SUBTEND	Submit end time (HHMMSS)
16	96	1	MSGTZONE	Time and date zone
17	97	1	MSGRCPTS	Receipts requested indicator (<i>H, T, blank</i>)
18	98	1	EXPAND	Expansion indicator
19	99	1	DTBLTYP	Alias table type (<i>G, O, P</i>)
20	100	3	DTBLID	Alias table name
21	103	1	SUBSTAD	Submit start date century (C)
22	104	1	SUBCEND	Submit end date century (C)

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

DESTACCT (field 5)

See "Understanding an Information Exchange address specification" on page 121.

DESTUID (field 6)

The format is alphanumeric, left-justified, and padded on the right with blanks. See “Understanding an Information Exchange address specification” on page 121.

DESTTYPE (field 7)

This field contains the value that determines the type of reference in DESTACCT (field 5) and DESTUID (field 6).

This Value: Indicates:

D A local true ID or an alias reference.

L The name of a distribution list.

MSGCLASS (field 8)

This field contains a value that identifies the class of delivery service for messages you want to cancel.

This value: Indicates:

blank Normal priority.

P High priority.

MSGNAME (field 9)

If this field is not blank, it must match the message name on the first or only message of any group you want to cancel. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGSEQN (field 10)

If this field is not blank, it must match the input message sequence number on the first or only message of any group you want to cancel. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGUCLS (field 11)

If this field is not blank, it must match the user classification field on the first or only message of any group that you want to cancel. The format is alphanumeric, left-justified, and padded on the right with blanks.

SUBDSTA (field 12)

This field contains a start date that specifies a time range in which the submit time of the first or only message of a group you want to cancel must fall, offset to the time zone specified in MSGTZONE (field 16). If blank, the value defaults to *000102*. The format is numeric (*YYMMDD*), right-justified, and padded on the left with zeros.

SUBTSTA (field 13)

This field contains a time that specifies the start time of a time range in which the submit time of the first or only message of a group you want to cancel must fall, offset to the time zone specified in MSGTZONE (field 16). If blank, the value defaults to 000000. The format is numeric (*HHMMSS*), right-justified, and padded on the left with zeros.

SUBDEND (field 14)

This field contains a date that specifies the end date of a time range in which the submit time of the first or only message of a group you want to cancel must fall, offset to the time zone specified in MSGTZONE (field 16). If blank, this field defaults to 420916 for expansion level 2, and 991231 for expansion level 1. The format is numeric (*YYMMDD*), right-justified, and padded on the left with zeros.

SUBTEND (field 15)

This field contains a time that specifies the end time of a time range in which the submit time of the first or only message of a group you want to cancel must fall, offset to the time zone specified in MSGTZONE (field 16). If blank, this field defaults to 235959. The format is numeric (*HHMMSS*), right-justified, and padded on the left with zeros.

MSGTZONE (field 16)

This field contains the reference time zone for SUBDSTA (field 12), SUBTSTA (field 13), SUBDEND (field 14), and SUBTEND (field 15).

This value:	Indicates:
L	The date and time are offset to the time zone specified in the Session Start command.
G	The date and time are offset to Greenwich mean time.
blank	This value defaults to G .

MSGRCPTS (field 17)

This field contains a value requesting that Information Exchange send an acknowledgment message regarding the cancellation of a message group to your user ID. There are two types of acknowledgments; both are receipt acknowledgments that Information Exchange creates and queues to your user ID each time it cancels a message group. The first type contains only the message header information about the first or only message in the canceled message group. The second type contains the header information and the message text of the first or only message in the canceled message group.

This value:	Indicates:
blank	No receipt acknowledgment is created.
H	Only header information is included in the acknowledgments.
T	Both header and text information are included in the acknowledgments.

EXPAND (field 18)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

DTBLTYP (field 19)

The format is alphanumeric. See “Understanding an Information Exchange address specification” on page 121.

DTBLID (field 20)

The format is alphanumeric, left-justified, and padded on the right with blanks. See “Understanding an Information Exchange address specification” on page 121.

SUBCSTA (field 21)

This field contains the start date century when expansion level 2 is requested for the Cancel command. A value of “0” indicates a 19xx year; a value of “1” indicates a 20xx year. The default value is 0.

SUBCEND (field 22)

This field contains the end date century when expansion level 2 is requested for the Cancel command. A value of “0” indicates a 19xx year; a value of “1” indicates a 20xx year. The default value is 1.

[Understanding an Information Exchange address specification](#)

The value in DESTTYPE (field 7) indicates the address specification used. A value of *D* in this field indicates a local true ID or an alias reference. A value of *L* indicates a distribution list.

A local true ID is specified as follows:

Field:	Contains:
DESTTYPE (field 7)	<i>D</i>
DTBLTYP (field 19)	blank
DTBLID (field 20)	blank
DESTACCT (field 5)	Destination account ID
DESTUID (field 6)	Destination user ID

An alias reference is specified as follows:

Field:	Contains:
DESTTYPE (field 7)	<i>D</i>
DTBLTYP (field 19)	Table type (<i>G</i> , <i>O</i> , or <i>P</i>)
DTBLID (field 20)	Table name
DESTACCT (field 5)	First 8 characters of the alias name
DESTUID (field 6)	Last 8 characters of the alias name

A distribution list is specified as follows:

Field:	Contains:
DESTTYPE (field 7)	<i>L</i>
DTBLTYP (field 19)	blank
DTBLID (field 20)	blank
DESTACCT (field 5)	Distribution list name
DESTUID (field 6)	blank

Date processing for expansion level 1

For processing of commands issued in the 20th century:

- The first 2 characters of the year for the start and end date are 19.

For processing of commands issued in the 21st century:

If the year portion of the start date field (SUBDSTA) is greater than 42, then the date will default as follows:

- For the start date, the first 2 characters of the year are 19.
- If the year portion of the end date field (SUBDEND) is greater than 42, then the first 2 characters of the year for the end date are 19.
- If the year portion of the end date field (SUBDEND) is less than or equal to 42, then the first 2 characters of the year for the end date are 20.

If the year portion of the start date field (SUBDSTA) is less than or equal to 42, then the date will default as follows:

- For the start date, the first 2 characters of the year are 20.
- For the end date, the first 2 characters of the year are 20.

Commit command—SDICMIT

This command causes Information Exchange to pause for synchronization, while receiving messages from your Information Exchange address. It also establishes a session-restart point for submitting messages. This command causes Information Exchange to checkpoint the messages received from an Information Exchange address; its length is 48 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDICMIT
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	RETRSPC	Response command name
6	41	5	MSGCNTR	Message segments count
7	46	2	COMITCD	Commit code
8	48	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

RETRSPC (field 5)

This field contains the value that is returned in the first 8 characters of the Commit response.

MSGCNTR (field 6)

This field contains the number of message segments, exclusive of ENDDATA messages, your system sent to Information Exchange since the last Commit or Pace command (or since the beginning of the session, if there has been no prior Commit or Pace command). Information Exchange does not test the value, but reflects its own count of message segments received in the Commit response. The format is numeric, right-justified, and padded on the left with zeros.

COMITCD (field 7)

This field contains *00*. Any value other than zero in this field causes Information Exchange to end the session abnormally. The session access key is invalidated. In order to resume communication with Information Exchange, a new Session Start command should be issued. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 8)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

Commit command response

This is Information Exchange's response to the Commit command; its length is 45 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDICMIT
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	5	CHKPNTNO	Checkpoint number
6	38	2	CHKPNTCD	Checkpoint response code
7	40	5	CHKPNTMC	Number of message segments received
8	45	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

CHKPNTNO (field 5)

This field contains the checkpoint number that Information Exchange assigns. Use this number in SNDCKPTN (field 14) of the Session Start command to tell the system to restart the session from the last checkpoint. The format is numeric, right-justified, and padded on the left with zeros.

CHKPNTCD (field 6)

This field contains the Commit response code. If this value is other than zero, it indicates that Information Exchange has not assumed responsibility for all messages sent since the last Commit command and response exchange. The value here is the highest severity code Information Exchange detects in message-header editing and message routing since the last Commit or Pace exchange. The format is numeric, right-justified, and padded on the left with zeros.

CHKPNTMC (field 7)

This field contains the number of message segments, exclusive of ENDDATA messages, received by Information Exchange since the previous Commit or Pace command (or since the start of the session). To validate that this is the correct number of message segments sent since the last Commit or Pace exchange, you can:

- Send a Session End command coded to end the session abnormally, permitting Information Exchange to ignore these messages
- Send any other input that causes Information Exchange to process the checkpointed messages to their destinations by Information Exchange.

The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 8)

This field contains the value from EXPAND (field 8) of the Commit command.

Pace command—SDIPACE

This command causes Information Exchange to pause for synchronization, while receiving messages from your Information Exchange address. It also enables your system to test the number of message segments sent since the last Commit or Pace command, or since the beginning of the session; its length is 48 characters.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIPACE
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	RETRSPC	Response command name
6	41	5	MSGCNTR	Message segments transferred
7	46	2	PACECOD	Pace code
8	48	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

RETRSPC (field 5)

This field contains the value that is returned in the first 8 characters of the Pace response. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGCNTR (field 6)

This field contains the number of message segments, exclusive of ENDDATA messages, sent to Information Exchange since the last Pace or Commit command, or since the beginning of the session (if there has been no prior Pace or Commit command). Information Exchange does not test the value but reflects its own count of message segments received in the Pace response. The format is numeric, right-justified, and padded on the left with zeros.

PACECOD (field 7)

This field contains a Pace code. Placing a value other than zero in this field causes Information Exchange to end the session abnormally. The session access key is invalidated. In order to resume communication with Information Exchange, a new Session Start command should be issued. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 8)

This field contains a command expansion-level indicator. It must contain the value *1*. Invalid values default to blank.

Pace command response

This is Information Exchange's response to the Pace command; its length is 45 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIPACE
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	5	reserved	Reserved. Must be blank.
6	38	2	PACECOD	Pace code
7	40	5	PACERCMC	Number of message segments received
8	45	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

RESERVED (field 5)

This field is reserved; leave it blank.

PACECOD (field 6)

This field contains the value *00* or a value indicating errors in the message segments preceding the last Pace or Commit command. The value here is the highest severity code Information Exchange detects during message header editing and message routing since the last Commit or Pace exchange. The format is numeric, right-justified, and padded on the left with zeros.

PACERCMC (field 7)

This field contains the number of message segments, exclusive of ENDDATA messages, received by Information Exchange since the previous Pace or Commit command, or since the start of the session. You might want to validate that this is the number of messages since the last Pace or Commit exchange. You can then:

- Send a Session End command that is coded to cause the session to end abnormally
- Send any other input to cause the session to remain active.

The format is numeric, right-justified, and padded on the left with blanks.

EXPAND (field 8)

This field contains the value from EXPAND (field 8) of the previous Pace command.

Probe command—SDIPROB

With this command, you can predetermine the validity of Information Exchange addresses to which you want to send messages, and the message payment levels associated with those Information Exchange addresses. The length of this command is 63 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIPROB
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	1	CMDMODE	Command mode and response placement (A, R, C, L)
8	43	1	MSGCHRG	Message service charges
9	44	20	DESTID	Destination ID

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Probe response. The format is alphanumeric, left-justified, and padded on the right with blanks.

CMDMODE (field 7)

This field contains a value that indicates the command mode and response placement. This field determines whether Information Exchange will send you an Asynchronous or Synchronous Probe response.

This value:	Indicates:
A (ASYNCHRONOUS)	The command is to be asynchronous; the response is placed in your Information Exchange mailbox for later retrieval.
R (RESPONSE)	The command is to be synchronous and the response is sent immediately.
C (COMPLETE)	The command is to be synchronous and the response is sent immediately.
L (LIST)	The command is to be synchronous, and a response is sent immediately. The command is issued against a distribution list, instead of against a user. In this case, DESTID (field 9) contains the list name.

If this field contains any other value, the system assumes a value of *C* and processes the command accordingly without generating error messages.

MSGCHRG (field 8)

This field contains the same value that would be placed in MSGCHRG (field 17) of a Send Message command to validate the charge scheme. If anything other than a numeric character from 1 through 6 is placed in this field, Information Exchange assumes a value of 3 and processes the command accordingly, without generating error messages.

DESTID (field 9)

This field contains the destination identifier. When *A*, *R*, or *C* is specified, the following applies.

- True ID

This value:	Indicates:
1 to 4	Blank.
4 to 12	The account ID.
13 to 20	The user ID.

- Alias reference

This value:	Indicates:
1	The alias table type (G, O, or P).
2 to 4	The alias table ID.
5 to 20	The alias name.

- Intersystem ID

This value:	Indicates:
1	I.
2 to 4	The intersystem system ID.
5 to 12	The account ID.
13 to 20	The user ID.

Probes to intersystem destinations ADV or IMX result in a RESPCODE value of 1. Probes to other intersystem destinations result in a RESPCODE value of 4.

When *L* is specified.

This value:	Indicates:
1 to 8	The list name.
9 to 20	Ignore.

Asynchronous Probe response message text

This is Information Exchange's response to the Probe command if an asynchronous response was requested; it is placed in your mailbox as a single Information Exchange message from an account ID of *SYSTEM* and a user ID of *PRBRSP*. The length of this response is 21 bytes. See "Received Message header" on page 169 for specific fields in the header.

#	COL	SIZE	NAME	DESCRIPTION
1	1	20	DESTID	Destination ID
2	21	1	RESPCODE	Destination validity indicator

DESTID (field 1)

This field contains the destination ID as sent on the Probe command.

RESPCODE (field 2)

This field contains a value that indicates the type of response.

This value: Indicates:

- | | |
|---|--|
| 1 | Messages can be sent to the indicated destination. All attempted probes to intersystem destinations of ADV or IMX return this code. If the probe is directed to a distribution list, the response indicates that the list exists, but does not imply the ability to send to any particular user in the list. Use the List Verify command to determine which users in a list are valid and which are invalid. |
| 2 | The destination is invalid. If the probe is directed to a list, this response indicates that the list does not exist. |
| 3 | The destination is valid; however, the trading partner's list prevents you from sending to it. |
| 4 | Information Exchange is unable to verify this destination. All attempted probes to intersystem destinations, other than ADV and IMX, return this code. |

Output format

The Asynchronous Probe response builds a message group in your Information Exchange mailbox. The following are in the message group's identification fields:

This field:	Contains:
DESTACCT	*SYSTEM*
DESTUID	*PRBRSP*
MSGNAME	blanks
MSGSEQN	blanks
MSGNCLS	blank
MSGUCLS	#SPROBE
MSGCLASS	blanks
MSGRCPTS	blank
MSGCHRG	5
SYSTYPE	IBMIE
SYSLEVEL	Current Information Exchange Version and Release level
DTBLTYP	blank
DTBLID	blanks

Synchronous Probe response

This is Information Exchange's response to the Probe command if a synchronous response was requested; its length is 56 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIPROB
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	ERRCODE	Error code
7	36	20	DESTID	Destination ID
8	56	1	RESPCODE	Destination validity indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the value from EXPAND (field 5) of the Probe command.

ERRCODE (field 6)

This field contains error code values.

This value: Indicates:

00 There was no error processing the command.

04 There was a recoverable error. You specified an invalid value and the default was used.

DESTID (field 7)

This field contains the destination as specified in the Probe command.

RESPCODE (field 8)

This field contains a value that indicates the type of response.

This value:	Indicates:
1	Messages can be sent to the indicated destination. All attempted probes to intersystem destinations of ADV or IMX return this code. If the probe is directed to a distribution list, the response indicates that the list exists, but does not imply the ability to send to any particular user in the list. Use the List Verify command to determine which users in a list are valid and which are invalid.
2	The destination is invalid. If the probe is directed to a list, this response indicates that the list does not exist.
3	The destination is valid; however, the trading partner's list prevents you from sending to it.
4	Information Exchange is unable to verify this destination. All attempted probes to intersystem destinations, other than ADV and IMX, return this code.

Purge Input command—SDIPURG

This command is used when the message sending program recognizes there is an error in the message, such as a mismatch between the ISA and IEA for X.12 messages. The sending program then issues the SDIPURG command, which causes the last uncommitted sent message group to be deleted. The length of the command is 33 bytes. There is no response returned to the submitter of the command.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIPURG
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a command expansion-level indicator. It must contain the value 1.

Send Message command—SDISNDM

With this command, you can send messages or message groups to trading partners. The long format is used for the first (or only) segment of a message; its length is 125 bytes, plus the length of the message or segment. The short format is used for the second and subsequent segments of multiple-segment messages; its length is 39 bytes, plus the length of the segment.

Long format

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDISNDM
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	MSGTXTOV	Continued message indicator (blank, <i>S</i>)
6	34	5	MSGTXTLN	Message text length
7	39	1	MSGGRPIN	Messages grouping indicator (<i>S, C, L, E</i>)
8	40	8	DESTACCT	Destination account ID or distribution list or first 8 characters of destination alias name
9	48	8	DESTUID	Destination user ID or last 8 characters of destination alias name
10	56	1	DESTTYPE	Destination type (<i>D, L</i>)
11	57	8	MSGNAME	Message name
12	65	5	MSGSEQN	Message sequence number
13	70	1	MSGNCLS	Network message classification (blank, <i>T, S, R</i>)
14	71	8	MSGUCLS	User message classification
15	79	1	MSGCLASS	Message delivery class (blank, <i>P, I</i>)
16	80	1	MSGRCPTS	Receipt requested indicator (blank, <i>R, D, B, A, C, E, F</i>)
17	81	1	MSGCHRG	Message service charges
18	82	2	MSGRETN	Undelivered retention time
19	84	1	SELRCV	Force receive search criteria (<i>F</i> , blank)
20	85	15	reserved	Reserved. Must be blank.
21	100	1	EXPAND	Expansion indicator
22	101	8		Reserved. Must be blank.
23	109	13	RESERVED	Reserved. Must be blank.
24	122	1	DTBLTYP	Alias table type (<i>G, O, P</i>) or intersystem indicator (<i>I</i>)
25	123	3	DTBLID	Alias table name or intersystem system ID
26	126	var	MSGTEXT	User's message text

Short format

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDISNDM
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	MSGTXTOV	Continued message indicator (<i>C, E</i>)
6	34	5	MSGTXTLN	Message text length
7	39	1	EXPAND	Expansion indicator
8	40	var	MSGTEXT	User's message text

COMMAND (field 1)**ACCNTNO (field 2)****USERID (field 3)****SESSKEY (field 4)****MSGTXTOV (field 5)**

The value in this field identifies the type of header (long or short) and indicates message-segment continuation.

This value:	Indicates:
blank	This is the only segment of a single-segment message; it uses the long-format header.
S	This is the first segment of a multiple-segment message; it uses the long-format header.
C	This is an intermediate segment of a multiple-segment message; it uses the short-format header.
E	This is the last segment of a multiple-segment message; it uses the short-format header.

A valid sequence of message segments for MSGTXTOV (field 5) is as follows:

... B ... (only)
 ... S ... (first)
 ... C ... (second)
 ... E ... (third)
 ... S ... (first)
 ... E ... (second)

MSGTXTLN (field 6)

This field contains a value that indicates the length of the message text that immediately follows this header. You can leave this field blank, in which case Information Exchange inserts the message-text size when it receives the message. This field can also be a numeric value, in which case Information Exchange checks it against the length of the text it received. The format is numeric, right-justified, and padded on the left with zeros.

MSGGRP (field 7)

This field contains a value that identifies message grouping or an end-of-data message.

This value: Indicates:

blank	A single message, not part of a group of messages.
S	The first message of a group, where Information Exchange delivers all messages within the group as a package, with no intervening messages from any other source.
C	A subsequent message within a group.
L	The last message of a group.
E	An ENDDATA message. In this case, there is no text, and the text length field, MSGTXTLN (field 6), contains zeros or blanks.

DESTACCT (field 8)

The format is alphanumeric, left-justified, and padded on the right with blanks. See “Understanding destination address specification” on page 141.

DESTUID (field 9)

The format is alphanumeric, left-justified, and padded on the right with blanks. See “Understanding destination address specification” on page 141.

DESTTYPE (field 10)

This field contains an address specification:

This value: Indicates:

D	True local ID, intersystem ID, or alias reference.
L	The name of a distribution list.

See “Understanding destination address specification” on page 141.

MSGNAME (field 11)

This field contains the name of a message or group of messages. If the message is the first in a group, this field names the entire group. This field is not validated by Information Exchange. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGSEQN (field 12)

This field may be used by the submitter to assign a message-control number to a message, or to number messages within a message group. This field is not validated by Information Exchange. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGNCLS (field 13)

This network message class field determines certain special process classifications for a message.

This value: Indicates:

blank	A normal message with no CDH.
T	A test-mode message with no CDH.
S	A normal message with CDH available.
R	A test-mode message with CDH available.

MSGUCLS (field 14)

This field enables you to indicate the type of data in a message. You can use it to indicate a printer form number or any other general classification agreed upon by the end users of Information Exchange. Information Exchange does not reference this field. However, addressees may reference it in the Receive Message command to select which messages on their queue they want to receive. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGCLASS (field 15)

The value in this field identifies the class of delivery service for this message. For a message group, the message class of the first message within the group determines the delivery-service class for the entire group.

This value: Indicates:

blank	A normal-priority delivery.
P	A high-priority delivery.
I	An express-priority delivery. Express-priority messages are delivered only if the destinations are currently receiving messages from Information Exchange. Express-priority message groups are not delivered until all messages within the group are received by Information Exchange, and a subsequent recovery point is reached.

MSGRCPTS (field 16)

This field contains a value that specifies the type of acknowledgment messages you want to receive. Acknowledgments returned to cluster mailboxes are sent to the cluster child user ID that issued the Send command.

This value:	Indicates:
blank	Information Exchange does not create acknowledgments.
R	Information Exchange creates only a receipt acknowledgment.
D	Information Exchange creates only delivery acknowledgments.
B	Information Exchange creates both receipt and delivery acknowledgments.
A	Information Exchange creates only a purge acknowledgment.
C	Information Exchange creates both a receipt and a purge acknowledgment.
E	Information Exchange creates either a purge or a delivery acknowledgment.
F	Information Exchange creates a receipt acknowledgment, and either a purge or a delivery acknowledgment.

For more information, refer to the discussion of Information Exchange acknowledgments in *Information Exchange Messages and Formats*.

MSGCHRG (field 17)

This field contains a value that indicates to Information Exchange how the sender wants the message charges to be paid as detailed in the table below. For a message group, the charge method requested in the first message within the group determines the charge mode for the entire group. The values for this field are the following:

Message charge code:	Means:
1	The receiver pays all the charges for the messages.
2	The receiver pays all the charges. If the receiver does not agree to pay all the charges, the sender and the receiver split the charges.
3	The receiver pays all the charges. If the receiver does not agree to pay all the charges, the sender and the receiver split the charges. If the receiver does not agree to split the charges, the sender pays all the charges.
4	The sender and the receiver split the charges. If the receiver does not agree to split the charges, the sender pays all the charges.
5	The sender and the receiver split all the charges.
6	The sender pays all the charges.

MSGRETN (field 18)

Use this field to specify the number of days that Information Exchange keeps the message group stored if it is not received by the destination. This field is effective only for the first (or only) long message header of a message group. If this field does not contain valid numeric digits, or if 00 is specified, the system supplies a default value. The format is numeric digits. The maximum value that may be specified is system defined. Information Exchange ignores any value greater than the maximum and uses the default.



NOTE: The field size will remain at 2 digits and the user is responsible for encoding the number of days as follows: The high-order digit may be 0-9 or A-I. If A-I is used, it is translated to mean 10-18. The low-order digit may be 0-9 and will remain translated as 0-9. That is, the user would supply a value of F6 to indicate a message retention of 156 days.

SELRCV (field 19)

This field indicates whether or not the recipient of this message must specify selective receive search criteria for the message to be delivered. Valid values are *F* and **blank**. If *F* is specified, one of the following receive search criteria must be used:

- Sender's address
- User message class
- Message key

The default is **blank**, indicating the message can be received without specifying the selective receive search criteria.

RESERVED (field 20)

This field is reserved; leave it blank.

EXPAND (field 21—long format, or field 7—short format)

This field contains a command expansion-level indicator. It must contain the value *1*. Invalid values default to blank.

RESERVED (field 22)

This field is reserved; leave it blank.

RESERVED (field 23)

This field is reserved; leave it blank.

DTBLTYP (field 24)

See “Understanding destination address specification” on page 141.

DTBLID (field 25)

This field is alphanumeric, left-justified, and padded on the right with blanks. For more information on alias tables, See “Understanding destination address specification” on page 141.

MSGTEXT (field 26—long format, or field 8—short format)

This field contains your message text.

Understanding destination address specification

The value in DESTTYPE (field 8) indicates the address specification used. A value of D in this field indicates a local true ID, an intersystem ID, or an alias reference. A value of L indicates a distribution list. A local true ID is specified as follows:

Field:	Contains:
DESTTYPE (field 10)	<i>D</i>
DTBLTYP (field 24)	blank
DTBLID (field 25)	blank
DESTACCT (field 8)	Destination account ID
DESTUID (field 9)	Destination user ID

An intersystem ID is specified as follows:

Field:	Contains:
DESTTYPE (field 10)	<i>D</i>
DTBLTYP (field 24)	<i>I</i>
DTBLID (field 25)	Intersystem system ID
DESTACCT (field 8)	Destination account ID
DESTUID (field 9)	Destination user ID

An alias reference is specified as follows:

Field:	Contains:
DESTTYPE (field 10)	<i>D</i>
DTBLTYP (field 24)	Table type (<i>G, O, or P</i>)
DTBLID (field 25)	Table name
DESTACCT (field 8)	First 8 characters of the alias name
DESTUID (field 9)	Last 8 characters of the alias name

A distribution list is specified as follows:

Field:	Contains:
DESTTYPE (field 10)	<i>L</i>
DTBLTYP (field 24)	blank
DESTBLID (field 25)	blank
DESTACCT (field 8)	Distribution list name
DESTUID (field 9)	blank

Mailbox commands

Commit request

This request causes a checkpoint synchronization between your system and the messages Information Exchange sends to your user ID; its length is 48 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from COMMCMD of SDISSTA
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	CMTRSPC	Response command name
6	41	5	MSGCNTR	Message segments count
7	46	2	COMITCD	Commit code
8	48	1	EXPAND	Expansion indicator

COMMAND (field 1)

COMMCMD (field 12) of the Session Start command.

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

CMTRSPC (field 5)

This field contains the value SDICRSP. It must be placed in the first 8 characters of the Commit response returned to Information Exchange. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGCNTR (field 6)

This field contains the number of message segments, exclusive of ENDDATA messages, Information Exchange sent to your user ID since the last Commit request (or since the beginning of the session, if there has been no prior Commit request). The format is numeric, right-justified, and padded on the left with zeros.

COMITCD (field 7)

This field contains 00. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 8)

This field contains the value from EXPAND (field 17) of your Session Start command.

Commit request response—SDICRSP

You send this response to a Commit request from Information Exchange; its length is 45 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDICRSP
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	5	CHKPNTNO	Receive checkpoint number
6	38	2	CHKPNTCD	Checkpoint response code
7	40	5	CHKPNTMC	Message segments transferred
8	45	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

CHKPNTNO (field 5)

This field contains the checkpoint number your system assigns. It must be a number other than zero, and it must differ from prior checkpoint numbers sent in this session. The format is numeric, right-justified, and padded on the left with zeros.

CHKPNTCD (field 6)

This field contains a Commit response code. Placing a value here other than zero causes Information Exchange to end the session abnormally. The session access key is invalidated. In order to resume communication with Information Exchange, a new Session Start command should be issued. The format is numeric, right-justified, and padded on the left with zeros.

CHKPNTMC (field 7)

Information Exchange does not process this field when it receives a Commit response. It is your system's responsibility to determine whether message counts disagree. If the counts disagree, the CHKPNTCD field can then be set to a value other than zero to cause Information Exchange to end a session abnormally. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 8)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

Load Test Messages command—SDILTST

This command identifies test messages you want to receive and causes them to be added to your normal-priority message queue. You can identify test messages from an account ID of *SYSTEM* and a user ID of *TSTMSG*. The length of this command is 39 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILTST
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	3	LOMSGNO	Low message number
6	36	3	HIMSGNO	High message number
7	39	1	EXPAND	Expansion indicator

COMMAND (field 1)**ACCNTNO (field 2)****USERID (field 3)****SESSKEY (field 4)****LOMSGNO (field 5)**

This field contains the number of the first of a series of test messages you requested. If LOMSGNO is the same as HIMSGNO (field 6), then Information Exchange sends only one message. The format is numeric, right-justified, and padded on the left with zeros.

HIMSGNO (field 6)

This field contains the number of the last of a series of test messages that you request. The field HIMSGNO must be greater than or equal to LOMSGNO (field 5). If HIMSGNO is the same as LOMSGNO, then Information Exchange sends only one message. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 7)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

Output format

The Load Test Messages command builds a message group in your Information Exchange mailbox. The following are in the message group's identification fields:

This field:	Contains:
DESTACCT	*SYSTEM*
DESTUID	*TSTMSG*
MSGNAME	blanks
MSGSEQN	blanks
MSGNCLS	blank
MSGUCLS	blanks
MSGCLASS	blanks
MSGRCPTS	blank
MSGCHRG	5
SYSTYPE	IBMIE
SYSLEVEL	Current Information Exchange Version and Release level
DTBLTYP	blank
DTBLID	blanks

The contents of the messages will contain one of the following:

If you specify message #	You receive:
0	Numerics 0 to 9
1	Uppercase alphabet
2	Lowercase alphabet
3	Numerics 0 to 9 and both uppercase and lowercase alphabet
4	Special characters, represented in hexadecimal form as follows: (‘4F7C7B5B6C5F505C4D5D606D7E4E5A4AE06AC0D07D7F5E7 A616F4B6B40’X)
5	All 256 EBCDIC characters in ascending sequence

Message Inquiry command—SDIINQM

This command tells Information Exchange to return information about messages currently available for retrieval from your mailbox. The length of this command is 41 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIINQM
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	RETRSPC	Inquiry response command code
6	41	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

RETRSPC (field 5)

This field contains the value that Information Exchange places in the first 8 characters of the Message Inquiry response. The format is alphanumeric, left-justified, and padded on the right with blanks.

EXPAND (field 6)

This field contains a command expansion-level indicator.

This value:	Indicates:
1 or blank	The Message Inquiry response will not include the CDH.
2	The Message Inquiry response will include the CDH length in DATAMSGS (field 7), SERVMSGs (field 9), and TOTLDATA (field 11).

Invalid values default to blank.

Message Inquiry response

This is Information Exchange's response to the Message Inquiry command; its length is 93 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIINQM
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	6	INQDATE	Date of inquiry
6	39	6	INQTIME	Time of inquiry
7	45	6	DATAMSGS	Data messages available
8	51	12	DATAAGE	Age of oldest data message
9	63	6	SERVMSGs	Service messages available
10	69	12	SERVAGE	Age of oldest service message
11	81	6	TOTLDATA	Total message data available
12	87	6	MAXSIZ	Size of largest message or group
13	93	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

INQDATE (field 5)

This field contains the date on which the inquiry was processed, offset to the time zone specified in the Session Start command. The format is numeric (*YYMMDD*) and right-justified.

INQTIME (field 6)

This field contains the time at which the inquiry was processed, offset to the time zone specified in the Session Start command. The format is numeric (*HHMMSS*) and right-justified.

DATAMSGS (field 7)

This field contains the number of user message data segments available. A message data segment is measured in terms of the data unit size specified in *MAXMSGsZ* (field 10) of the Session Start command. The value is calculated using the total size of all user messages minus the size of all message header data. The format is numeric, right-justified, and padded on the left with zeros.

DATAAGE (field 8)

This field contains the age of the oldest data message available. The format is numeric (*YYMMDDHHMMSS*—your local time) and right-justified.

SERVMSGs (field 9)

This field contains the number of Information Exchange message data segments available. A message data segment is measured in terms of the data unit size specified in *MAXMSGsZ* (field 10) of the Session Start command. The value is calculated using the total size of all Information Exchange messages minus the size of all message header data. The format is numeric, right-justified, and padded on the left with zeros.

SERVAGE (field 10)

This field contains the age of the oldest Information Exchange message available. The format is numeric (*YYMMDDHHMMSS*—your local time), and right-justified.

TOTLDATA (field 11)

This field contains a value that indicates the total number of user and Information Exchange message data segments available. The format is numeric, right-justified, and padded on the left with zeros.

MAXSIZ (field 12)

This field contains the amount of storage required to store the largest single message or message group available. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 13)

This field contains the value from EXPAND (field 6) of the Message Inquiry command.

Message Queue Query command—SDIQUMS

You can use this command to browse the contents of your Information Exchange mailbox. You can then decide what action you want to take on these messages. The length of this command is 62 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIQUMS
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	1	TYPRSPN	Response type requested (<i>1 or 3</i> —no CDH, <i>2 or 4</i> —CDH)
8	43	20	LOCATOR	Command chaining value

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a command expansion-level indicator. It must contain the value 1, 2, or 3. Invalid values default to blank. An expansion level 1 does not return a receive indicator, an archive indicator, or a no delete indicator in the message queue entry. An expansion level 2 returns a receive indicator in the message queue entry. An expansion level 3 returns a receive indicator, an archive indicator, and a no delete indicator in the message query entry.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Message Queue Query response. The format is alphanumeric, left-justified, and padded on the right with blanks.

TYPRSPN (field 7)

This field contains a value that determines the type of response that Information Exchange generates.

This value:	Indicates:
1	Only the Information Exchange header data is shown in the response. Messages in the process of being received are not returned.
2	The Information Exchange header data plus the CDH data for those message groups with common data headers are shown in the response. Messages in the process of being received are not returned.
3	Only the Information Exchange header data is shown in the response. Messages in the process of being received are returned.
4	The Information Exchange header data plus the CDH data for those message groups with common data headers are shown in the response. Messages in the process of being received are returned.

LOCATOR (field 8)

This field contains a command chaining value, which the system uses to receive information for additional message queue values. To use this field:

- Issue a Message Queue Query command with 20 zeros in this field.
- Issue Message Queue Query commands with LOCATOR set to CHAINRSP (field 8) of the previous Message Queue Query response, until CHAINRSP contains a series of Fs.

If this field contains a series of Fs, there are no more messages in the queue.

Message Queue Query response

This is Information Exchange's response to the Message Queue Query command; its length is 60 bytes, plus the amount of data in the message queue entries.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIQUMS
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	ERRCODE	Error code
7	36	1	TYPRSPN	Response type requested
8	37	20	CHAINRSP	Command chaining value
9	57	4	ITEMS	Number of message queue items
10	61	n	MSG QUEUE ENTRY	First message queue entry

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator sent on the Message Queue Query command.

ERRCODE (field 6)

This field specifies error code values.

This value:	Indicates:
00	There was no error processing the command.
04	There was a recoverable error; one or more items may be present.
08	There was a nonrecoverable error; no items are present.

If ERRCODE is not 00, issue a Receive Message command to retrieve the associated error messages.

TYPRSPN (field 7)

This field contains the value from TYPRSPN (field 7) of the Message Queue Query command.

CHAINRSP (field 8)

This field is a command chaining value to be copied into the LOCATOR field of subsequent Message Queue Query commands. If this value contains a series of Fs, there are no more message groups that meet your specifications.

ITEMS (field 9)

This field contains the number of message queue entries that follow.

MSG QUEUE ENTRY (field 10)

This field contains the first message queue entry (see “Message Queue entry” on page 153.)

There can be more than one instance of this field, as indicated by the value in ITEMS (field 9).

Message Queue entry

Expansion level 1

This is the format of a single message queue entry; its length is 96 bytes plus the CDH if applicable.

#	COL	SIZE	NAME	DESCRIPTION
1	1	20	MSGID	Message ID string
2	21	20	MSGSNDR	Message sender
3	41	8	MSGNAME	Message name
4	49	5	MSGSEQ	Message input sequence
5	54	1	MSGNCLAS	Message network class
6	55	8	MSGUCLS	User message class
7	63	8	SYSTYPE	Sender's system type
8	71	4	SYSLEVEL	Sender's system level
9	75	10	MSGTXTSZ	Size of message group text
10	85	6	MSGDATE	Date message available
11	91	6	MSGTIME	Time message available
12	97	n	common data header	Common data header (CDH) data

Expansion level 2

This is the format of a single message queue entry; its length is 97 bytes plus the CDH if applicable.

#	COL	SIZE	NAME	DESCRIPTION
1	1	20	MSGID	Message ID string
2	21	20	MSGSNDR	Message sender
3	41	8	MSGNAME	Message name
4	49	5	MSGSEQ	Message input sequence
5	54	1	MSGNCLAS	Message network class
6	55	8	MSGUCLS	User message class
7	63	8	SYSTYPE	Sender's system type
8	71	4	SYSLEVEL	Sender's system level
9	75	10	MSGTXTSZ	Size of message group text
10	85	6	MSGDATE	Date message available
11	91	6	MSGTIME	Time message available
12	97	1	RCVIND	Receive indicator
13	98	n	common data header	Common data header (CDH) data

Expansion level 3

This is the format of a single message queue entry; its length is 99 bytes plus the CDH if applicable.

#	COL	SIZE	NAME	DESCRIPTION
1	1	20	MSGID	Message ID string
2	21	20	MSGSNDR	Message sender
3	41	8	MSGNAME	Message name
4	49	5	MSGSEQ	Message input sequence
5	54	1	MSGNCLAS	Message network class
6	55	8	MSGUCLS	User message class
7	63	8	SYSTYPE	Sender's system type
8	71	4	SYSLEVEL	Sender's system level
9	75	10	MSGTXTSZ	Size of message group text
10	85	6	MSGDATE	Date message available
11	91	6	MSGTIME	Time message available
12	97	1	RCVIND	Receive indicator
13	98	1	ARCHIND	Archive indicator
14	99	1	NODELIND	No delete indicator
15	100	n	common data header	Common data header (CDH) data header

MSGID (field 1)

This field contains the message identifier that Information Exchange assigns. It is the hexadecimal representation of the message delivery class and a unique identifier.

MSGSNDR (field 2)

This field contains the sender's identification. If the sender is on a different system, the first character is an I and the following 3 characters are the system identifier; if the sender is on the same system that you are on, the first 4 characters are blank. This is followed by the account ID (8 characters) and the user ID (8 characters) of the message sender.

MSGNAME (field 3)

This field contains the message name that the message sender assigns. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGSEQ (field 4)

This field contains the message input sequence that the message sender assigns. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGNCLAS (field 5)

This field contains the message network class that the message sender and your trading partner list assigns.

This value:	Indicates:
blank	Normal message with no CDH available.
T	Test mode message with no CDH available.
S	Normal message with CDH available.
R	Test mode message with CDH available.



NOTE: If this field contains R or S and TYPRSPN (field 7) of the Message Queue Query command contains a 2, then the CDH for this message group immediately follows MSGTIME (field 11) below; otherwise, the system does not return CDH data.

MSGUCLS (field 6)

This field contains the user message class that the message sender assigns.

SYSTYPE (field 7)

This field contains the message sender's system type, as it is set in the Session Start command.

SYSLEVEL (field 8)

This field contains the message sender's system level, as it is set in the Session Start command.

MSGTXTSZ (field 9)

This field contains the total text size, exclusive of any Information Exchange headers or common data headers. The format is numeric.

MSGDATE (field 10)

This field contains the date the message is available for receipt, offset to the time zone specified in the Session Start command. The format is *YYMMDD*.

MSGTIME (field 11)

This field contains the time that the message is available for receipt, offset to the time zone specified in the Session Start command. The format is *HHMMSS*.

RCVIND (field 12, expansion levels 2 and 3 only)

This field contains the receive indicator.

This value:	Indicates:
Y	The message is in the process of being received.
F	The message was marked for force selective receive criteria.
blank	The message is not in the process of being received nor was it marked for force selective receive criteria.

ARCHIND (field 13, expansion level 3 only)

This field contains the archive indicator.

This value:	Indicates:
Y	The message was retrieved from archive.
F	The message cannot be archived.
blank	The message was not retrieved from archive.

NODELIND (field 14, expansion level 3 only)

This field contains the no delete indicator.

This value:	Indicates:
Y	The message cannot be deleted. Some installations do not allow the deletion of specific types of messages.
blank	The message may be deleted.

Common Data Header (CDH) data (field 12, expansion level 1; field 13, expansion level 2; field 15, expansion level 3)

If `typrspn` (field 7 of the Message Queue Query command) contains a 2 or a 4 and the `MSGNCLAS` (field 5) contains an R or S, the CDH is included in the response.

Information Exchange provides a 4-byte numeric character field followed by the CDH. The CDH always begins with a 2-byte length field that describes the length of the entire header (including the length field itself) and is followed by a variable number of entries. For more information on CDHs, see “Common data header (CDH)” on page 23.

Pace request

This is a request from Information Exchange. It causes a pause for synchronization during the transmission of messages from Information Exchange to you; its length is 48 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from COMMCMD of SDISSTA
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	RETRSPC	Response command name
6	41	5	MSGCNTR	Message segment count
7	46	2	PACECOD	Pace code
8	48	1	EXPAND	Expansion indicator

COMMAND (field 1)

This field contains the value from commcmd of the Session Start command.

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

RETRSPC (field 5)

This field contains SDIPRSP, which is placed in the first 8 characters of the Pace request response returned to Information Exchange. The format is left-justified and padded on the right with blanks.

MSGCNTR (field 6)

This field contains the number of message segments, exclusive of ENDDATA messages, Information Exchange has sent to your user ID since the last Pace request, or since the beginning of the session (if there has been no prior Pace request). The format is numeric, right-justified, and padded on the left with zeros.

PACECOD (field 7)

This field contains a Pace response code. Placing a value other than zero in this field causes Information Exchange to end the session abnormally. The session access key is invalidated. In order to resume communication with Information Exchange, a new session start command should be issued. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 8)

This field contains the value from EXPAND (field 17) of the Session Start command.

Pace request response—SDIPRSP

This is the response that you send to Information Exchange as a result of Information Exchange's Pace request; its length is 45 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIPRSP
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	5	reserved	Reserved. Must be blank.
6	38	2	PACECOD	Pace response code
7	40	5	PACERCMC	Message segment count
8	45	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

RESERVED (field 5)

This field is reserved; leave it blank.

PACECOD (field 6)

This field contains a Pace response code. Placing a value other than zero in this field causes Information Exchange to end the session abnormally. The session access key is invalidated. In order to resume communication with Information Exchange, a new Session Start command should be issued. The format is numeric, right-justified, and padded on the left with zeros.

PACERCMC (field 7)

Information Exchange does not process this field when it receives a Pace response. You must determine if message counts disagree. If the counts disagree, you can set the PACECOD field to a value other than zero to cause an abnormal session end. The format is numeric, right-justified, and padded on the left with zeros.

EXPAND (field 8)

This field contains a command expansion-level indicator. It must contain the value *1*. Invalid values default to blank.

Purge Message command—SDIPRGM

With this command, you can delete messages from your Information Exchange mailbox without receiving them. To authorize use of this command, the “Use purge message command” flag in your Information Exchange profile must be set to Y. Refer to *Information Exchange Administration Services User’s Guide* for more information. The length of the command is 61 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIPRGM
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	20	MSGKEY	Message key to delete

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a command expansion-level indicator. It must contain the value 1. This value defaults to 1 if EXPAND (field 5) of the Purge Message command does not contain a 1.

RETRSPC (field 6)

This field contains the value that is returned in the first eight characters of the purge message response.

MSGKEY (field 7)

This field contains the 20-character message identifier that is returned in a response to a previously processed Message Queue Query command. The value in the MSGKEY field identifies the message that is to be purged from your Information Exchange mailbox.

Purge Message response

This is Information Exchange's response to a Purge Message command. The length of this response is 35 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIPRGM
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RESPONSE	Response code

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the value from EXPAND (field 5) of the purge message command.

RESPONSE (field 6)

This field contains error code values.

This value:	Indicates:
00	Normal completion; the message was purged from your mailbox.
01	The input command is too short to contain all required fields.
02	The message to be purged does not exist.
03	The message is being received and, therefore, cannot be purged.
04	The message cannot be purged because your Information Exchange user profile indicates that use of the Purge Message command is not authorized.
05	The message cannot be purged. Some installations do not allow the purging of specific types of messages.



NOTE: If any invalid values are detected in any field for which a default value is defined, a warning message is placed in the user's mailbox to notify the user that the default value was assumed. This may occur regardless of the value returned in the "response" field.

Receive Message command—SDIRCVM

You can use this command to receive messages that are in your mailbox. You can receive all messages, or you can receive messages meeting specific criteria. The selection criteria can be message key or one or more of the following:

- Sender's address or messages from archive
- User message class
- EDI message indicator
- Sent date and time

If a message is sent with the force receive search criteria indicator, then one of the following criteria must be used:

- Sender's address
- User message class
- Message key

If the expansion level is equal to 1, then the length of this command is 116 bytes and contains fields 1-21. If the expansion level is equal to 2, then this command is 136 bytes and contains the fields 1-22. If the expansion level is equal to 4 or 5, then this command is 163 bytes and contains the fields 1-27.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIRCVM
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	REFNAME	Command reference name
6	41	8	DESTACCT	Source account ID or distribution list or first 8 characters of source alias name
7	49	8	DESTUID	Source user ID or last 8 characters of source alias name
8	57	1	DESTTYPE	Source type (blank, <i>D</i> , <i>L</i> , <i>A</i>)
9	58	8	MSGUCLS	User message classification
10	66	8	QMSGTRID	Stored message command
11	74	8	IMSGTRID	Express delivery message command
12	82	8	SMSGTRID	Information Exchange message command
13	90	1	TYPECMND	Type of SDIRCVM command (blank, <i>C</i> , <i>E</i> , <i>G</i> , <i>W</i>)
14	91	5	MAXMSGNO	Maximum number of messages
15	96	3	reserved	Reserved
16	99	1	EDIRCV	Receive EDI messages indicator (blank, <i>Y</i> , <i>N</i>)
17	100	4	reserved	Reserved
18	104	1	EXPAND	Expansion indicator
19	105	1	DTBLTYP	Alias table type (<i>G</i> , <i>O</i> , <i>P</i>) or intersystem indicator (<i>I</i>)

#	COL	SIZE	NAME	DESCRIPTION
20	106	3	DTBLID	Alias table name or intersystem system ID
21	109	8	ARCREFID	Archive reference ID
22	117	20	MSGKEY	Message group key
23	137	7	DATEFROM	From date used for messages selection (CYMMDD)
24	144	6	TIMEFROM	From time used for message selection (HHMMSS)
25	150	7	DATETO	To date used for message selection (CYMMDD)
26	157	6	TIMETO	To time used for message selection (HHMMSS)
27	163	1	TIMEZONE	Time zone for the DATEFROM, TIMEFROM, DATETO, TIMETO fields

COMMAND (field 1)**ACCNTNO (field 2)****USERID (field 3)****SESSKEY (field 4)****REFNAME (field 5)**

This field contains an identifier that distinguishes this command from others that might be simultaneously active. The format is alphanumeric, left-justified, and padded on the right with blanks.



NOTE: If you submit a subsequent Receive Message command that contains the same value in this field as an active, previously submitted Receive Message command, Information Exchange modifies the original command.

DESTACCT (field 6)

The format is alphanumeric, left-justified, and padded on the right with blanks. See “Understanding source address specification” on page 168.

DESTUID (field 7)

The format is alphanumeric, left-justified, and padded on the right with blanks. See “Understanding source address specification” on page 168.

DESTTYPE (field 8)

The default is blank.

This value: Indicates:

blank	The command is to receive any messages queued to your user ID, regardless of the originating user ID.
D	The command is to receive any messages queued to your user ID from a specific originating user ID.
L	The command is to receive any messages queued to your user ID from a list of originating user IDs.
A	The command is to receive all archived messages that have been requeued to your normal-priority message queue.

For more information, see “Understanding source address specification” on page 168.

MSGUCLS (field 9)

This field contains a user classification. You can limit the messages received to those with the user classification specified here. Information Exchange sends only those messages or message groups whose originators specified the same value in MSGUCLS (field 14) of the Send Message commands. You can use a question mark (?) as a *wild card* to substitute for any character or characters. For example, AB1? selects all users whose user classification begins with AB1; ?999 selects all those that end with 999. The format is alphanumeric, left-justified, and padded on the right with blanks.

QMSGTRID (field 10)

This field contains the value that the system places in the first 8 characters of the Received Message header for non-system, non-express delivery messages. The format is alphanumeric, left-justified, and padded on the right with blanks.

IMSGTRID (field 11)

This field contains the value that the system places in the first 8 characters of the Received Message header for any express delivery messages. The format is alphanumeric, left-justified, and padded on the right with blanks.

SMSGTRID (field 12)

This field contains the value that the system places in the first 8 characters of the Received Message header for system-generated messages. The format is alphanumeric, left-justified, and padded on the right with blanks.

TYPECMND (field 13)

This field contains a value that establishes the type of receive operation to be performed.

This value:	Indicates:
blank	A batch-receive operation. When Information Exchange has sent all data queued for your user ID, Information Exchange also sends an ENDDATA message (with the value from field 10 placed in field 1 of the Receive Message Header—long format).
C	A continuous-receive operation. Information Exchange sends data as it is available. Information Exchange does not send an ENDDATA message until the value specified in MAXMSGNO (field 14) is reached or the continuous-receive is ended by issuing another Receive Message command with the same value in REFNAME (field 5) and a TYPECMND (field 13) value of <i>E</i> .
E	Information Exchange must end the prior continuous-receive operation. Information Exchange continues to send data until any incomplete message or message group has been transmitted. Your system must be prepared to accept message segments from Information Exchange until ENDDATA is received.
G	A single-group receive operation. Information Exchange sends one message group and the ENDDATA message.
W	A wait-to-receive operation, which is a combination of the continuous receive and the single message group receive. A single message group is received, and the receive process ends. However, if there is no message in the mailbox that currently qualifies according to the selection criteria, the system does not end the receive, but places the session in a “wait” status. The system waits for the arrival of a qualified message.

MAXMSGNO (field 14)

If TYPECMND (field 13) contains a value of blank, *C*, or *W*, this field contains the maximum number of message segments that you are prepared to receive. If you do not want transmission to end, specify *00000*; otherwise, the system sends the ENDDATA message following the number of message segments specified in this field.

If TYPECMND contains the value *G*, this field contains the maximum number of message segments you are prepared to receive for the group. If you specify *00000*, Information Exchange sends the message group, regardless of size; otherwise, it transmits the message group only if the group can be delivered with this number of message segments or fewer.

The format is numeric, right-justified, and padded on the left with zeros.

RESERVED (field 15)

This is a reserved field; leave it blank.

EDIRCV (field 16)

This field contains a value that indicates whether you want to receive EDI messages, non-EDI messages, or both. Specify *Y* in this field to receive only EDI messages. Specify *N* to receive only non-EDI messages. If you want to receive both types of messages, leave this field blank.



NOTE: Only those messages identified in the CDH as having one of the EDI formats in the DFORMAT field are sent to your system if this field contains a *Y*.

RESERVED (field 17)

This is a reserved field; leave it blank.

EXPAND (field 18)

This field contains a command expansion-level indicator.

This value:	Indicates:
blank and 1	No CDH is received, and MSGKEY (field 22) is not present in the command.
2	CDH is received, and MSGKEY (field 22) is present in the command.
4	CDH is received, and fields 22 through 27 are present in the command.
5	CDH is received, and fields 22 through 27 are present in the command. MSGSIZE (field 28) is returned in the long format received message header.

DTBLTYP (field 19)

See “Understanding source address specification” on page 168.

DTBLID (field 20)

The format is alphanumeric, left-justified, and padded on the right with blanks. See “Understanding source address specification” on page 168.

ARCREFID (field 21)

This field contains a value that indicates that all messages received by this command are to be placed in the short-term archive. The value in this field identifies messages to be retrieved by the Archive Retrieve command. If this field is left blank and the user’s profile indicates that archive is to be forced, the session access key is used as an archive reference.



NOTE: The combination of a relatively short archive retention period and a relatively long continuous-receive operation might cause you to lose messages intended for archiving. For example, if you define the short-term archive retention period as three days, issue a Receive Message command that specifies continuous-receive (*C* in TYPECMND) and leave the command in effect for longer than three days. Those messages received after the third day are archived, but they are immediately eligible for purging. See “Archiving messages” on page 59 for more information.

MSGKEY (field 22)

To receive a specific message group, insert the message group key as obtained with a Message Queue Query command in this field. Information Exchange sends that single message group followed by an ENDDATE message to your system. This operates as though *G* were specified in TYPECMND (field 13), regardless of the value there.



NOTE: Field 22 is present only if EXPAND (field 18) is 2 or 4.

DATEFROM (field 23)

This field is used when receiving messages sent to the mailbox within a specified time frame. To receive messages sent to the mailbox within a specified time frame, you can specify the start date for the search. Information Exchange will extract only messages sent on or after the date you specified in this field. Enter the date in the format *CYYMMDD*, where *C* indicates the century, *YY* indicates the low-order digits of the year, *MM* indicates the month, and *DD* indicates the day of the month. The following centuries are recognized:

This value:	Indicates:
0	1900
1	2000

The default value for this field is 0000102 (January 2, 1900). If this field is blank, the default value is applied.



NOTE: Field 23 is present only if EXPAND (field 18) is 4.

TIMEFROM (field 24)

This field is used when receiving messages sent to the mailbox within a specified time frame. To receive messages sent to the mailbox within a specified time frame, you can specify the start time for the search. Information Exchange will extract only messages sent on or after the time you specified in this field. Enter the time in the format *HHMMSS*, where *HH* indicates the hour, *MM* indicates the minutes, and *SS* indicates the seconds. The field is right-justified and padded on the left with zeros. The default value for this field is 000000 (00:00:00 AM). If this field is blank, the default value is applied.



NOTE: Field 24 is present only if EXPAND (field 18) is 4.

DATETO (field 25)

This field is used when receiving messages sent to the mailbox within a specified time frame. To receive messages sent to the mailbox within a specified time frame, you can specify the end date for the search. Information Exchange will extract only messages sent on or prior to the date you specified in this field.

Enter the date in the format *CYYMMDD*, where *C* indicates the century, *YY* indicates the low-order digits of the year, *MM* indicates the month, and *DD* indicates the day of the month. The following centuries are recognized:

This value:	Indicates:
0	1900
1	2000

The default value for this field is 1420916 (September 16, 2042). If this field is blank, the default value is applied.



NOTE: Field 25 is present only if EXPAND (field 18) is 4.

TIMETO (field 26)

This field is used when receiving messages sent to the mailbox within a specified time frame. To receive messages sent to the mailbox within a specified time frame, you can specify the end time for the search. Information Exchange will extract only messages sent on or prior to the time you specified in this field. Enter the time in the format *HHMMSS*, where *HH* indicates the hour, *MM* indicates the minutes, and *SS* indicates the seconds. The field is right-justified and padded on the left with zeros. The default value for this field is 235959 (11:59:59 PM). If this field is blank, the default value is applied.



NOTE: Field 26 is present only if EXPAND (field 18) is 4.

TIMEZONE (field 27)

This field is used when receiving messages sent to the mailbox within a specified time frame. To receive messages sent to the mailbox within a specified time frame, you can specify the time zone to be applied for the DATEFROM, TIMEFROM, DATETO, TIMETO fields. The following values are recognized:

This value:	Indicates:
L	The time zone specified in the Session Start command
G	Greenwich mean time

The default value for this field is L (local time). If this field is blank, the default value is applied.



NOTE: Field 27 is present only if EXPAND (field 18) is 4.

Understanding source address specification

The value in DESTTYPE (field 8) indicates the address specification used. A value of D in this field indicates a local true ID, an intersystem ID, or an alias reference. A value of L indicates a distribution list.

A local true ID is specified as follows:

Field:	Contains:
DESTTYPE (field 8)	<i>D</i>
DTBLTYP (field 19)	blank
DTBLID (field 20)	blank
DESTACCT (field 6)	Source account ID
DESTUID (field 7)	Source user ID

An intersystem ID is specified as follows:

Field:	Contains:
DESTTYPE (field 8)	<i>D</i>
DTBLTYP (field 19)	<i>I</i>
DTBLID (field 20)	Intersystem system ID
DESTACCT (field 6)	Source account ID
DESTUID (field 7)	Source user ID

An alias reference is specified as follows:

Field:	Contains:
DESTTYPE (field 8)	<i>D</i>
DTBLTYP (field 19)	Table type (<i>G, O, or P</i>)
DESTBLID (field 20)	Table name
DESTACCT (field 6)	First 8 characters of the alias name
DESTUID (field 7)	Last 8 characters of the alias name

A distribution list is specified as follows:

Field:	Contains:
DESTTYPE (field 8)	<i>L</i>
DTBLTYP (field 19)	blank
DTBLID (field 20)	blank
DESTACCT (field 6)	Distribution list name
DESTUID (field 7)	blank

Received Message header

A message header identified the message to your system. The long format is present for the first (or only) segment of a message. For expansion levels 1, 2, and 4, the length of this header is 125 bytes, plus the length of the message or message segment. For expansion level 5, the length of the header is 135, plus the length of the message or message segment. The short-format message header is present in the second and subsequent segments of a multiple-segment message. The length of this header is 39 bytes, plus the length of the segment.

Long format -- Expansion levels 1, 2, and 4

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from QMSGTRID, IMSGTRID, or SMSGTRID of SDIRCVM
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	MSGTXTOV	Continued message indicator
6	34	5	MSGXTLN	Message text length
7	39	1	MSGGRPIN	Message grouping indicator
8	40	8	DESTACCT	Source account ID
9	48	8	DESTUID	Source user ID
10	56	1	DESTTYPE	Destination type
11	57	8	MSGNAME	Message or group name
12	65	5	MSGSEQN	Message sequence number
13	70	1	MSGNCLS	Network message classification
14	71	8	MSGUCLS	User message classification
15	79	1	MSGCLASS	Message delivery class
16	80	1	MSGRCPTS	Receipts requested indicator
17	81	1	MSGCHRG	Message service charge indicator
18	82	6	MSGSEQO	Message sequence out number
19	88	8	SYSTYPE	Source system type
20	96	4	SYSLEVEL	Source system level
21	100	1	EXPAND	Expansion indicator
22	101	8	EXTRTN	Reserved
23	109	6	MSGDATE	Message date
24	115	6	MSGTIME	Message time
25	121	1	MSGTZONE	Time and date zone
26	122	1	DTBLTYP	Alias table type (<i>G</i> , <i>O</i> , <i>P</i>) or intersystem indicator (<i>I</i>)
27	123	3	DTBLID	Alias table name or intersystem system ID
28	126	var	MSGTEXT	User's message text

Long format -- Expansion level 5

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from QMSGTRID, MSGTRID, or SMSGTRID of SDIRCVM
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	MSGTXTOV	Continued message indicator
6	34	5	MSGTXTLN	Message text length
7	39	1	MSGGRPIN	Message grouping indicator
8	40	8	DESTACCT	Source account ID
9	48	8	DESTUID	Source user ID
10	56	1	DESTTYPE	Destination type
11	57	8	MSGNAME	Message or group name
12	65	5	MSGSEQN	Message sequence number
13	70	1	MSGNCLS	Network message classification
14	71	8	MSGUCLS	User message classification
15	79	1	MSGCLASS	Message delivery class
16	80	1	MSGRCPTS	Receipts requested indicator
17	81	1	MSGCHRG	Message service charge indicator
18	82	6	MSGSEQO	Message sequence out number
19	88	8	SYSTYPE	Source system type
20	96	4	SYSLEVEL	Source system level
21	100	1	EXPAND	Expansion indicator
22	101	8	EXTRTN	Reserved
23	109	6	MSGDATE	Message date
24	115	6	MSGTIME	Message time
25	121	1	MSGTZONE	Time and date zone
26	122	1	DTBLTYP	Alias table type (<i>G</i> , <i>O</i> , <i>P</i>) or intersystem indicator (<i>I</i>)
27	123	3	DTBLID	Alias table name or intersystem system ID
28	126	10	MSGSIZE	Length of message text without CDH
29	136	var	MSGTEXT	User's message text

Short format

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from QMSGTRID, IMSGTRID, or SMSGTRID of SDIRCVM
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	MSGTXTOV	Continued message indicator
6	34	5	MSGTXTLN	Message text length
7	39	1	EXPAND	Expansion indicator
8	40	var	MSGTEXT	User's message text

COMMAND (field 1)

This field contains the value from QMSGTRID (field 10), IMSGTRID (field 11), or SMSGTRID (field 12) of the Receive Message command.

This type of message:	Value from:
Non-system, non-express delivery	QMSGTRID
Express delivery	IMSGTRID
System-generated	SMSGTRID

ACCNTNO (field 2)**USERID (field 3)****SESSKEY (field 4)****MSGTXTOV (field 5)**

This field contains a message overflow indicator.

This value:	Indicates:
blank	A single-segment message. This is the only segment of the current message; it uses the long-format header.
S	The first segment of a multiple-segment message; it uses the long-format header.
C	An intermediate segment of a multiple-segment message; it uses the short-format header.
E	The last segment of a multiple-segment message; it uses the short-format header.

MSGTXTLN (field 6)

This field contains the length of MSGTEXT (field 28) in this segment. The format is numeric, right-justified, and padded on the left with zeros.

MSGGRPIN (field 7)

This field contains a message-grouping indicator.

This value: Indicates:

blank	A single message and not part of a group of messages.
S	The first message of a group, where all messages within the group are delivered as a package, with no intervening messages from any other source.
C	A continuation message within a group.
L	The last message of a group.
E	An ENDDATA message with no text and a MSGTXTLN (field 6) containing zeros. Information Exchange sends an ENDDATA message when any of the following is true: <ul style="list-style-type: none"> • The maximum number of messages you are prepared to receive (indicated by MAXMSGNO (field 14) of the Receive Message command) has been sent. • All of the selected messages you requested using a noncontinuous Receive command have been sent. • You end a continuous-receive command.



NOTE: If there are messages meeting the specified selection criteria remaining in the mailbox after the completion of this receive command, the characters MORE, left-justified and padded on the right with blanks, is placed in MSGNAME (field 11 of the long format header). The following conditions must be true in order for the characters MORE to be assigned:

TYPECMND (field 13 of the Receive Message command) is NOT C.

MAXMSGNO (field 14 of the Receive Message command) is NOT zeros or blank or the system limit on message groups not committed is exceeded.

EXPAND (field 18 of the Receive Message command) is 4.

MSGGRPIN (field 7 of the long format header) is E.

DESTACCT (field 8)

This field, in conjunction with DESTUID (field 9), identifies the message sender. If the value in this field is *SYSTEM*, the message was generated by Information Exchange. (See “Messages generated by Information Exchange” on page 57.) The format is alphanumeric, left-justified, and padded on the right with blanks.

DESTUID (field 9)

This field, in conjunction with DESTACCT (field 8), identifies the message sender. The format is alphanumeric, left-justified, and padded on the right with blanks.

DESTITUTE (field 10)

This field contains the value *D*, indicating the source identified by DESTACCT (field 8) and DESTUID (field 9) is a single user ID.

MSGNAME (field 11)

This field contains the message name (or group name) that the message sender assigns. It is not validated by Information Exchange, but the message sender can use it to give a name to a message or to a group of messages. The format is alphanumeric, left-justified, and padded on the right with blanks.



NOTE: If there are messages remaining in the mailbox meeting the specified selection criteria after the completion of this receive command, the characters MORE, left-justified and padded on the right with blanks, is placed in this field. The following conditions must be true for the characters MORE to be assigned:

- TYPECMND (field 13 of the Receive Message command) is **NOT** C.
- MAXMSGNO (field 14 of the Receive Message command) is **NOT** zeros or blank or the system limit on message groups not committed is exceeded.
- EXPAND (field 18 of the Receive Message command) is 4.
- MSGGRPIN (field 7 of the long-format header) is E.

MSGSEQN (field 12)

This field contains a message sequence number that the message sender assigns. This field is not validated by Information Exchange. The format is numeric, right-justified, and padded on the left with zeros.

MSGNCLS (field 13)

This network message class field determines certain special process classifications for this message.

This value:	Indicates:
blank	A normal message without a CDH.
T	A test-mode message without a CDH.
S	A normal message with CDH available.
R	A test-mode message with CDH available.

MSGUCLS (field 14)

This field contains the value of the message user class assigned to the message by the sender.

MSGCLASS (field 15)

This field identifies the delivery class of this message. The format is alphanumeric, left-justified, and padded on the right with blanks.

This value:	Indicates:
blank	A normal-priority delivery.
P	A high-priority delivery.
I	An express-priority delivery.
A	A normal-priority, requeued archive message.

MSGRCPTS (field 16)

This field contains either a blank, indicating that no delivery acknowledgment was sent to the message sender, or the value *D*, indicating that a delivery acknowledgment was sent to the message sender.

MSGCHRG (field 17)

This field contains a value that indicates how charges for the message are assigned.

This value: Indicates:

blank	There are no charges for this message.
1	That you, the receiver, are charged for both send- and receive-side charges. This is a collect message.
5	That you, the receiver, are charged for the receive-side charges, and the sender is charged for the send-side charges (split charges).
6	That the sender is charged for both the receive- and send-side charges (prepaid charges).

MSGSEQO (field 18)

This field contains a unique number (0 to 999999 and repeats, if necessary) that Information Exchange assigns to each message sent to a user ID. The format is numeric, right-justified, and padded on the left with zeros.

SYSTYPE (field 19)

This field contains the value of SYSTYPE (field 7) that was set by the sender's Session Start command. The format is alphanumeric, left-justified, and padded on the right with blanks.

SYSLEVEL (field 20)

This field contains the value of SYSLEVEL (field 8) that was set by the sender's Session Start command. The format is alphanumeric, left-justified, and padded on the right with blanks.

EXPAND (field 21—long format, or field 7—short format)

This field contains the value from EXPAND (field 18) of the Receive Message command.

EXTRTN (field 22)

This is a reserved field; leave it blank.

MSGDATE (field 23)

This field contains the date that the sender entered the message into the Information Exchange system, offset to the time zone specified in MSGTZONE (field 25). The format is numeric (*YYMMDD*), right-justified, and padded on the left with zeros.

MSGTIME (field 24)

This field contains the time that the sender entered the message into the Information Exchange system, offset to the time zone specified in MSGTZONE (field 25). The format is numeric (*HHMMSS*), right-justified, and padded on the left with zeros.

MSGTZONE (field 25)

This field indicates the offset used for the date and time in MSGDATE (field 23) and MSGTIME (field 24). If the value is *L*, the time zone specified in the Session Start command is used. If the value is *G*, the time zone specified is GMT.

DTBLTYP (field 26)

This field, if DESTTYPE (field 10) is *D*, identifies the table type of an alias table. Combined with DTBLID (field 27), this field uniquely defines the alias table in which to resolve the alias name created by combining DESTACCT (field 8) and DESTUID (field 9). If an alias name was used, this field contains one of the following values:

This value:	Indicates:
G	A global alias table.
O	An organizational alias table.
P	A private alias table.

If the sender is on a different system, the field contains the value *I*. If neither of these is the case, the field is blank.

DTBLID (field 27)

This field, if DESTTYPE (field 10) is the value *D*, identifies the table name of an alias table. Combined with DTBLTYP (field 26), this field uniquely defines the alias table in which to resolve the alias name created by combining DESTACCT (field 8) and DESTUID (field 9). The format is alphanumeric, left-justified, and padded on the right with blanks. If an alias was used, this field contains the name of the alias table. If the sender is on a different system and DTBLTYP (field 26) value is *I*, this field contains the ID of the sender's system. If neither of these is the case, the field is blank.

MSGSIZE (field 28—long format, expansion level 5 only)

This field contains the size of the message without the CDH.

MSGTEXT (field 28—long format, expansion levels 1, 2, and 4, or field 29—long format, expansion level 5, or field 8—short format)

This field contains the text of the message.

Alias commands

Alias Inquiry command—SDIINQA

This command determines the true Information Exchange address associated with a given alias. If the user of the alias is on a different system, this is also identified. The length of the Alias Inquiry command is 61 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIINQA
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	1	TBLTYPE	Alias table type (<i>G, O, P</i>)
8	43	3	TBLNAME	Alias table name
9	46	16	ALIAS	Alias name

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a command expansion-level indicator. It can contain the values 1, 2, or blank.

This value: Indicates:

blank or 1 The 16-character Information Exchange account ID and User ID part of the address is returned.

2 An expanded user ID that includes the Information Exchange system ID is returned.

RETRSPC (field 6)

This field contains the value that returns in the first 8 characters of the Alias Inquiry response.

TBLTYPE (field 7)

This field indicates the type of alias table in which you are working. The format is a coded value.

This value: Indicates

- G *A global alias table.* A global table is accessible by any Information Exchange user.
- O *An organizational alias table.* An organizational table is accessible by any Information Exchange user in the same account.
- P *A private alias table.* A private table is available only to its owner.

TBLNAME (field 8)

This field contains the 3-character alias table name. The format is alphanumeric, left-justified, and padded on the right with blanks.

ALIAS (field 9)

This field contains the alias name to be resolved.

Alias Inquiry response

This is the Information Exchange response to the Alias Inquiry command; its length is 73 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIINQA
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	1	TBLTYPE	Alias table type (<i>G, O, P</i>)
7	35	3	TBLNAME	Alias table name
8	38	16	ALIAS	Alias name
9	54	20	TRUEID	True user ID

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a value from EXPAND (field 5) of the Alias Inquiry command.

TBLTYPE (field 6)

This field contains a value from the Alias Inquiry command.

TBLNAME (field 7)

This field contains a value from the Alias Inquiry command.

ALIAS (field 8)

This field contains the value from the Alias Inquiry command.

TRUEID (field 9)

If the alias or table is unknown, this field contains blanks. Otherwise, it contains the true user ID associated with the alias. If the first character is an I, the user is associated with a different system. In this case, the next 3 characters are the system's identifier.

If the first 4 characters are blanks, then the indicated user is on the same Information Exchange system that you are on. The fifth through twentieth characters are the account ID (8 characters) and user ID (8 characters) of the user.



NOTE: The above is the extended Information Exchange address format. If the EXPAND level is less than two, this field is 16 characters long and contains only the account ID and user ID portions of the addressee.

Define Alias command—SDIDALS

This command defines a new alias or redefines an existing alias. You put aliases in an alias table and use them in place of a user's Information Exchange address. The length of the Define Alias command is 43 bytes, plus 36 bytes for every alias to be defined.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIDALS
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	TBLTYPE	Alias table type (<i>G, O, P</i>)
6	34	3	TBLNAME	Alias table name
7	37	1	TYPECMND	Command type (<i>A, N, D, C, E</i>)
8	38	1	AUTHCOD	Updated authorization code (<i>G, A, P</i>)
9	39	1	EXPAND	Expansion indicator
10	40	4	ITEMS	Number of following entries
11	44	36n	ALIAS ENTRY	Alias table entry

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

TBLTYPE (field 5)

This field contains the type of alias table in which you are working. The format is a coded value.

This value: Indicates:

- G A global alias table. A global table is accessible by any Information Exchange user.
- O An account or organization alias table. An organization table is accessible by any Information Exchange user in the same account.
- P A private alias table. A private table is only accessible by its owner.

TBLNAME (field 6)

This field contains the 3-character alias table name. The format is alphanumeric, left-justified, and padded on the right with blanks.

TYPECMND (field 7)

This field indicates what is to be done to the alias table. The format is a coded value.

This value: Indicates:

A	The command is to add the entries to an existing alias table.
N	The command is to create a new alias table.
D	The command is to delete the entries from an existing alias table.
C	The command is to change the entries in an existing alias table.
E	The entire alias table is to be erased.

AUTHCOD (field 8)

This field indicates who can update the alias table. The format is a coded value.

This value: Indicates:

G	Any user can update this global table.
A	Any administrator in the account can update the table.
P	Only the owner of the table can update it.



NOTE: If you create a global table using global update authorization, users outside your account can add or delete table entries that pertain to their own account, but only network support personnel can erase the entire table. Other authorization codes for global tables permit the owner or the account service administrator to make all changes, including erasing the table.

For more information about aliases, refer to *Information Exchange Administration Services User's Guide*.

EXPAND (field 9)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

ITEMS (field 10)

This field contains the number of alias entries. The format is numeric, right-justified, and padded on the left with zeros. Specify the value *0000*, for an erase table command (E-type command).

ALIAS ENTRY (field 11)

This field contains the alias definitions. Each alias definition uses 36 characters. For a description, see “Alias entry” on page 181.

Alias entry

#	COL	SIZE	NAME	DESCRIPTION
1	1	16	ALIAS	Alias name
2	17	1	TBLTYPE	Alias table type (<i>G</i> , <i>O</i> , <i>P</i>) or intersystem indicator (<i>I</i>)
3	18	3	TBLNAME	Alias table name or intersystem system ID
4	21	16	USER	Alias name reference or true account and user ID

ALIAS (field 1)

This field specifies the alias name being defined in the table.

TBLTYPE (field 2)

This field specifies whether the alias is being defined to point to another alias. You can chain up to 5 aliases. Specify a *G*, *O*, or *P* to refer to an alias table type. Specify *I* for an intersystem indicator. Specify *blank* if you are giving the true user ID.

TBLNAME (field 3)

This field specifies the alias table name if the value for TBLTYPE (field 2) is *G*, *O*, or *P*. If the value for TBLTYPE (field 2) is *I*, this field specifies the system identifier.

USER (field 4)

This field specifies the Account and user ID, alias name reference, or true account and user ID.

Distribution List commands

List Define command—SDILSTD

This command defines, modifies, or deletes a distribution list; its length is 47 bytes, plus 20 bytes for each Information Exchange address on the distribution list.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILSTD
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	LSTNAME	Distribution list ID
6	41	1	TYPECMND	Type of SDILSTD command (<i>A, N, D, E</i>)
7	42	1	LISTTYPE	Type of list (<i>P, A, G</i> , blank)
8	43	1	EXPAND	Expansion indicator
9	44	4	LSTSIZE	Number of destination IDs
10	48	20n	DESTIDS	Destination ID entry

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

LSTNAME (field 5)

This field contains the name of the list. The format is alphanumeric, left-justified, and padded on the right with blanks

TYPECMND (field 6)

This field contains an indicator for the type of operation to be performed on the list.

This value:	Indicates:
A	Add entries contained in DESTIDS (field 10) to a list by the name specified in LSTNAME (field 5). If no list is known by that name, the results are the same as with the value <i>N</i> in this field.
N	Create a new list or replace a previous list with the same name.
D	Delete the destinations contained in DESTIDS (field 10) from the named list. If all names are deleted, results are the same as with the value <i>E</i> in this field.
E	Erase the entire list named in LSTNAME (field 5).

LISTTYPE (field 7)

This field contains an indicator for the type of list.

This value:	Indicates:
P	The list is a user permanent list that remains after the end of the current Information Exchange session.
A	The list is an account-level list. This also is permanent, but is addressable by all users within the account.
G	The list is an account-level list that includes user grouping (See “G type lists” on page 183.). It is a permanent list that is addressable by all users within the account and only contains users within the same account.
blank	The list is a user temporary list that lasts only the duration of the session.

G type lists

The *G* (or group) list specifies a *user group* within the account. User groups provide an alternative method of authorization for intra-account message routing. A user can belong to one or more groups, or even to no group. The rules for using these groups are:

- A maximum of 80 groups can be defined for an account.
- A user within the account can be a member of any number of the groups.
- Groups can intersect to enable communication to flow across the intersection.
- A group can intersect itself to enable communication to flow within the group.
- A user’s ability to communicate within the account is the logical ORing (adding together) of the communication capabilities of all the groups containing that user.

All users in a group can communicate with all users in another group if the groups intersect. For example, assume there are three groups defined within the account:

- group 1
- group 2
- group 3

The following users are members of group 1:

- user a
- user b
- user c
- user d
- user e

The following users are members of group 2:

- user j
- user k
- user l
- user m
- user n

The following users are members of group 3:

- user a
- user b
- user c
- user j
- user k
- user l

Next assume that group 1 and group 3 intersect and that group 2 and group 3 intersect. Users a, b, c, j, k, and l can communicate with all the other users. This is because they are in group 3, which is connected to both the other groups. However, users d, e, m, and n can communicate only with the users named in group 3.

If group intersections are increased so that group 1 intersects with itself, and group 2 intersects with itself, all members of group 1 can communicate with all other members of group 1. Similarly, all members of group 2 can communicate with all other members of group 2.

EXPAND (field 8)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

LSTSIZE (field 9)

This field contains the number of destination IDs that follow if TYPECMND (field 6) is *A*, *N*, or *D*. The format is alphanumeric, right-justified, and padded on the left with zeros.

DESTIDS (field 10)

This field contains a list of 20-character destination IDs. The format is alphanumeric, left-justified, and padded on the right with blanks. Each destination ID must occupy 20 positions. See “Destination ID entry” on page 185 for a complete description.

Destination ID entry

#	COL	SIZE	NAME	DESCRIPTION
1	1	1	TBLTYPE	Alias table type (<i>G</i> , <i>O</i> , <i>P</i>) or intersystem indicator (<i>I</i>)
2	2	3	TBLNAME	Alias table name or intersystem system ID
3	5	16	USER	Alias name reference of true account/user ID

TBLTYPE (field 1)

This field specifies whether the alias is being defined to point to another alias. You can chain up to 5 aliases. Specify a *G*, *O*, or *P* to refer to an alias table type. Specify *I* for an intersystem indicator. Specify *blank* if you are giving the true user ID.

TBLNAME (field 2)

If character 1 is *G*, *O*, or *P*, this field names the alias table to be used to search for the name given in characters 5 through 20. If character 1 is *I*, this is the system identifier. If character 1 is blank, these characters are not used.

USER (field 3)

If character 1 is *G*, *O*, or *P*, this field specifies the 16-character alias name of the desired destination. If character 1 is *blank* or *I*, this field specifies the Information Exchange account (characters 5 through 12) and user ID (characters 13 through 20) of the desired destination.

List Verify command—SDILSTV

With this command, you can verify a list for sending or receiving messages, or for doing both; its length is 43 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILSTV
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	8	LSTNAME	Distribution list ID
6	41	1	TYPECMND	Type of SDILSTV command (<i>L, D, C, A, R, S, B</i>)
7	42	1	MSGCHRG	Message service charges
8	43	1	EXPAND	Expansion indicator

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

LSTNAME (field 5)

This field contains the name of the list you want to verify. The format is alphanumeric, left-justified, and padded on the right with blanks.

TYPECMND (field 6)

This field contains a value that indicates how you want Information Exchange to verify the list named in LSTNAME (field 5).

This value:	Indicates:
L	This is a list of users in a permanent list. A message is built with each 20 characters of message text in the format of the Destination ID entry. See “Destination ID entry” on page 185.
D	This is a list of a user’s permanent list. A message is built with each 8 characters of message text containing a list name.
C	This is a list of users in an account or group list. A message is built with each 20 characters of message text in the format of the Destination ID entry. See “Destination ID entry” on page 185.
A	This is a list of account and group lists. A message is built with each 8 characters of message text containing a list name.
R	This list is to be verified for receiving messages from the user IDs in the list. The user IDs in the list must be known to Information Exchange and be authorized to send messages to your user ID.

This value:	Indicates:
S	This list is to be verified for sending messages to the user IDs in the list. The user IDs in the list must be valid and authorized to receive messages from your user ID.
B	Both the receive and send verifications are to be performed. This is the default.

MSGCHRG (field 7)

This field contains a value that indicates which level of message-charge validation you want Information Exchange to perform when the value in TYPECMND (field 6) is *R*, *S*, or *B*. A blank in this field indicates the verification is to be done only on the basis of valid user ID and communication authority. Information Exchange does not test for message-charge.

Other values of this field are shown in Figure on page 139. The send verification ensures that your user ID is authorized to send to other users on the list a message with the corresponding MSGCHRG value. The receive verification ensures that each of the other users is authorized to send a message to your user ID with the corresponding MSGCHRG value.

EXPAND (field 8)

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to blank.

Output format

When the TYPECMND (field 6) value specified is *L*, *D*, *C*, or *A* and there are no errors with the command, the List Verify command builds a message group in your Information Exchange mailbox. The sender of the error message is *SYSTEM**ERRMSG*. The following are in the message group's identification fields:

This field:	Contains:
DESTACCT	*SYSTEM*
DESTUID	*LSTRSP*
DESTTYPE	D
MSGNAME	MSGNAME can be one of the following: GLIST —For account lists or groups PLIST —For permanent user lists
MSGSEQN	blank
MSGNCLS	blank
MSGUCLS	MSGUCLS can be one of the following: LIST NAME —Specified if contents of lists are requested LISTS —Specified if list of lists is requested
MSGCLASS	blanks
MSGRCPTS	blank
MSGCHRG	1
SYSTYPE	IBMIE
SYSLEVEL	Current Information Exchange version and release level
DTBLTYP	blanks
DTBLID	blanks

If the command contains an error, a message from *SYSTEM**ERRMSG* is sent to the mailbox.

When the TYPECMND (field 6) value specified is *S*, *R*, or *B* and the verification contains no errors, then no messages are sent to your mailbox. An error message is sent to your mailbox for every invalid user ID on the list.

Library commands

Handling library command responses

The value in the RETCODE field of any library command response indicates whether Information Exchange encountered any problems while processing the command.

This value:	Indicates:
00	Information Exchange did not encounter any problems while processing the command.
04	Information Exchange encountered a recoverable error and returned one or more warning messages to your mailbox. However, the command was processed successfully. Issue a Receive Message command to retrieve the warning messages.

There are two ways Information Exchange can return information concerning any errors found while processing library commands. The expansion level used when the command is submitted specifies which method is used.

Expansion level 1

A value of 08 in the RETCODE field indicates Information Exchange encountered a non-recoverable error. The command was not processed. Information Exchange returned one or more messages to your mailbox indicating the specific error condition.

Expansion level 2

The value in the RETCODE field indicates the specific error that Information Exchange encountered while processing the command. Depending on the error, Information Exchange may place additional information about the error in your mailbox.

Browse Library Member command—SDILBRW

The Browse Library Member command enables you to retrieve partial text of a library member; its length is 149 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILBRW
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Return response name
7	42	8	OWNER	Owning account ID
8	50	8	LIBNAME	Library name
9	58	8	MEMBER	Member name
10	66	4	LENGTH	Length of text
11	70	80	LOCATOR	Locator field

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1 or 2. Invalid values default to 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Browse Library Member response. The format is alphanumeric, left-justified, and padded on the right with blanks.

OWNER (field 7)

This field identifies the owner of the library LIBNAME (field 8) specified. The default value is your account ID. The format is alphanumeric, left-justified, and padded on the right with blanks.

LIBNAME (field 8)

This field contains the name of the library that contains the member to be browsed. The format is alphanumeric, left-justified, and padded on the right with blanks.

MEMBER (field 9)

This field contains the name of the member you want to browse. The format is alphanumeric, left-justified, and padded on the right with blanks.

LENGTH (field 10)

This field contains the maximum length of text to return. The format is numeric, right-justified, and padded on the left with zeros.

LOCATOR (field 11)

This field contains a command chaining value. It is used to receive the next response from Information Exchange, which is from the last break point. To use this field:

1. Issue the initial SDILBRW command with all zeros in this field.
2. Issue subsequent SDILBRW commands with LOCATOR set to the NEXTLOC value of the previous SDILBRW response.

Browse Library Member response

This is Information Exchange's response to the Browse Library Member command; its length is 120 bytes, plus the length of the data.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDILBRW
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Return code
7	36	1	CONTINUE	Continuation indicator
8	37	80	NEXTLOC	Locator field
9	117	4	LENGTH	Length of text
10	121	n	DATA	Text of library member

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator from the Browse Library Member command.

RETCODE (field 6)

This field specifies error code values in expansion level 1. See “Handling library command responses” on page 189

This field specifies error code values in expansion level 2.

This value:	Indicates:
05	Information Exchange encountered a nonnumeric character in LENGTH (field 10). One or more warning messages are returned to the mailbox.
10	You do not have read-access for this library. No messages are returned to the mailbox.
11	Information Exchange encountered an invalid value in LOCATOR (field 11). No messages are returned to the mailbox.
16	Information Exchange encountered an invalid value in OWNER (field 7). No messages are returned to the mailbox.
18	Information Exchange encountered an invalid or unknown library name in LIBNAME (field 8). No messages are returned to the mailbox.
19	Information Exchange encountered an invalid or unknown library member name in MEMBER (field 9). No messages are returned to the mailbox.
30	The read-access list specified for this library does not exist. Contact your service administrator. No messages are returned to the mailbox.

CONTINUE (field 7)

This field contains information on whether the member text continues.

This value:	Indicates:
0	Information Exchange came to the end of the library member text and no more data follows.
1	There is more library member text data to follow.

NEXTLOC (field 8)

Information Exchange uses this field to submit a subsequent SDILBRW command that continues with the next member in the library by copying it into the LOCATOR field.

LENGTH (field 9)

This field contains the length of the data returned.

DATA (field 10)

This field contains the text of the member. The length is defined by LENGTH (field 9).

Define Library command—SDILDEF

This command creates an Information Exchange library and the data identifies the name and the authorization levels applicable; its length is 166 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILDEF
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Command response name
7	42	8	LIBNAME	Library name
8	50	8	OWNACCT	Owning account ID
9	58	8	OWNUSRID	Owning user ID
10	66	1	WRITEVL	Library write access level (<i>P, O, G, L</i>)
11	67	8	WRITELST	Library write access list
12	75	1	READLVL	Library read access level (<i>P, O, G, L</i>)
13	76	8	READLIST	Library read access list
14	84	1	SUPPACC	Customer Assistance access
15	85	1	ACTION	Action code (<i>A, R</i>)
16	86	1	SEARCH	Searchable indicator
17	87	1	OWNRESP	Owner responsible for retrieve
18	88	79	LIBDESC	Library description

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1 or 2. Invalid values default to 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Define Library response. The format is alphanumeric, left-justified, and padded on the right with blanks.

LIBNAME (field 7)

This field contains the name of the library you want to define. The format is alphanumeric, left-justified, and padded on the right with blanks.

OWNACCT (field 8)

This field contains the owning account ID. The format is alphanumeric, left-justified, and padded on the right with blanks.

OWNUSRID (field 9)

This field contains the owning user ID. The format is alphanumeric, left-justified, and padded on the right with blanks.

WRITELVL (field 10)

This field contains the authority level for write access to the library. The format is a coded value.

This value:	Indicates:
P	Only the owner has write access to this library.
O	Only users within the same account have write access to this library.
G	Any user has write access to this library.
L	Any user in the list defined by the WRITELST specification has access to this library.

WRITELST (field 11)

This field contains the name of a permanent distribution list that details the users who can update the library.

READLVL (field 12)

This field contains the authority level for read access to the library. The format is a coded value.

This value:	Indicates:
P	Only the owner can read this library.
O	Only users within the same account can read this library.
G	Any user can read this library.
L	Any user in the list defined by the READLIST specification can read this library.

READLIST (field 13)

This field contains the name of a permanent distribution list that details the users who have read access to the library.

SUPPACC (field 14)

This field contains a code that specifies whether Customer Care has access to the library.

This value:	Indicates:
Y	Network support personnel have permission to read the data in the library.
N	Network support personnel do not have permission to read the data in the library.

If you specify a value other than *Y* or *N*, the system assumes a *Y* value. The default is **Y**.

ACTION (field 15)

This field contains a code that specifies what action Information Exchange is to take.

This value: Indicates:

A The library is to be added if it does not already exist. This is the default action.

R The library is to be replaced.

SEARCH (field 16)

This field contains a code that specifies whether the library is searchable. If you want the library to be searchable, specify a **Y**. If you do not want the library to be searchable, specify an **N**. The default is **N**.

OWNRESP (field 17)

This field contains a code that indicates whether the owner of the library is responsible for the charges associated with library retrieve commands.

This value: Indicates:

Y The send-side charges associated with subsequent Library Retrieve commands are charged to the library owner.

N or blank The send-side charges associated with subsequent retrieve commands are charged to the account that issued the Library Retrieve command. This is the default.

LIBDESC (field 18)

This field contains a description of the library.

Define Library response

This is Information Exchange's response to the Define Library command; its length is 35 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDILDEF
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Command return code

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator from the Define Library command.

RETCODE (field 6)

This field contains error code values for expansion level 1. See "Handling library command responses" on page 189.

This value:	Indicates:
00	There were no errors while processing this command.
04	A warning message, such as an invalid expansion level, has been placed in the mailbox. However, the library is defined successfully.
08	An error has occurred during the processing of this command.

This field contains error code values for expansion level 2.

This value:	Indicates:
05	Information Exchange encountered an invalid command parameter. Warning messages are returned to the mailbox for any of the following fields:
Field name:	Indicates:
SUPPACC (field 14)	This field must contain a <i>Y</i> , <i>N</i> , or a blank.
OWNRESP (field 17)	This field must contain a <i>Y</i> , <i>N</i> , or a blank.
SEARCH (field 16)	This field must contain a <i>Y</i> or an <i>N</i> . The default is <i>N</i> .
ACTION (field 15)	This field must contain an <i>A</i> or an <i>R</i> . The default is <i>A</i> .

This value:	Indicates:
08	There is insufficient message size to start the session. Check the mailbox for an error message.
09	Information Exchange encountered an invalid command parameter. Check the mailbox for an error message for any of the following:
	Field name: Indicates:
	WRITELST (field 11) This field contains embedded blanks.
	READLST (field 13) This field contains embedded blanks.
	WRITELVL (field 10) This field must contain an <i>O, P, G, L</i> , or blank.
	READLVL (field 12) This field must contain an <i>O, P, G, L</i> , or blank.
10	You do not have access authority to define this library. No messages are returned to the mailbox.
16	Information Exchange encountered an invalid or unknown account ID in OWNACCT (field 8). No messages are turned to the mailbox.
	Field name: Indicates:
	OWNACCT (field 8) This field contains embedded blanks or the account specified does not exist.
17	Information Exchange encountered an invalid or unknown user ID in OWNUSRID (field 9). No messages are returned to the mailbox.
21	This library already exists and no replace option is specified in ACTION (field 15). No messages are returned to the mailbox.

Delete Library command—SDILDEL

This command deletes a library. If the Force Delete indicator is on, the system deletes all members of the library. The length of this command is 58 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILDEL
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Return response name
7	42	8	OWNER	Owning account ID
8	50	8	LIBNAME	Library name
9	58	1	OPTIONS	Delete option (blank, <i>F</i>)

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1 or 2. Invalid values default to 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Delete Library response. The format is alphanumeric, left-justified, and padded on the right with blanks.

OWNER (field 7)

This field contains the identification the owner of the library specified by LIBNAME (field 8). The default value is the account to which you belong. The format is alphanumeric, left-justified, and padded on the right with blanks.

LIBNAME (field 8)

This field contains the name of the library you want to delete. The format is alphanumeric, left-justified, and padded on the right with blanks.

OPTIONS (field 9)

This field contains the delete options that follow:

This value:	Indicates:
blank	The system should delete the library only if it contains no members. This is the default value.
F	The system deletes all members of the library. This is the Force Delete indicator.

Delete Library response

This is Information Exchange's response to the Delete Library command; its length is 35 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDILDEL
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Command return code

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator from the Delete Library command.

RETCODE (field 6)

This field specifies error code values for expansion level 1. See "Handling library command responses" on page 189.

This field specifies error code values for expansion level 2.

This value: Indicates:

- | | |
|----|--|
| 09 | Information Exchange encountered an invalid option code in OPTIONS (field 9). This field must contain a blank or an <i>F</i> . Check the mailbox for an error message. |
| 10 | You do not have the authority to delete the library specified in LIBNAME (field 8). Contact your service administrator. No messages are returned to the mailbox. |
| 16 | The account specified in OWNER (field 7) is invalid or does not exist. No messages are returned to the mailbox. |
| 18 | The library specified in LIBNAME (field 8) is invalid or does not exist. No messages are returned to the mailbox. |
| 22 | This library has members and you did not specify F in OPTIONS (field 9). No messages are returned to the mailbox. |

Delete Library Member command—SDILDLM

This command deletes a member from a library; its length is 65 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILDLM
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Return response name
7	42	8	OWNER	Owning account ID
8	50	8	LIBNAME	Library name
9	58	8	MEMBER	Member name

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1 or 2. Invalid values default to 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Delete Library Member response. The format is alphanumeric, left-justified, and padded on the right with blanks.

OWNER (field 7)

This field contains the identification the owner of the library specified by LIBNAME (field 8). The default value is the account to which you belong. The format is alphanumeric, left-justified, and padded on the right with blanks.

LIBNAME (field 8)

This field contains the name of the library from which you want to delete the specified member. The format is alphanumeric, left-justified, and padded on the right with blanks.

MEMBER (field 9)

This field contains the name of the member you want to delete. The format is alphanumeric, left-justified, and padded on the right with blanks.

Delete Library Member response

This is Information Exchange's response to the Delete Library Member command; its length is 35 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDILDLM
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Command return code

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator from the Delete Library Member command.

RETCODE (field 6)

This field specifies error code values for expansion level 1. See "Handling library command responses" on page 189.

This field specifies error code values for expansion level 2.

This value: Indicates:

- | | |
|----|--|
| 10 | You do not have write-access authority for the library specified in LIBNAME (field 8). No messages are returned to the mailbox. |
| 16 | You specified an invalid or unknown account ID in OWNER (field 7). No messages are returned to the mailbox. |
| 18 | You specified an invalid or unknown library name in LIBNAME (field 8). No messages are returned to the mailbox. |
| 19 | You specified an invalid or unknown member name in MEMBER (field 9). No messages are returned to the mailbox. |
| 30 | The write-access list for the library name in LIBNAME (field 8) does not exist. Contact your service administrator. No messages are returned to the mailbox. |

List Libraries command—SDILLST

This command returns a list of account libraries or a list of libraries for which you have either read or write access that satisfy the selection criteria. The length of this command is 84 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILLST
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Return response code
7	42	8	LIBACCT	Account to search
8	50	1	SELECT	Selection indicator (<i>A, C</i>)
9	51	1	AUTHSEL	Authority selection (<i>R, W</i>)
10	52	16	FORWLOCT	Forward continuation locator
11	68	16	BACKLOCT	Backward continuation locator
12	84	1	DIRINDC	Direction indicator (<i>F, B</i>)

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1 or 2. Invalid values default to 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the List Library response. The format is alphanumeric, left-justified, and padded on the right with blanks.

LIBACCT (field 7)

This field contains the account for which libraries are to be listed. If you leave this field blank, the system lists all libraries in the accounts you have access to.

SELECT (field 8)

This field contains the type of list required. The format is a coded value.

This value: Indicates:

A A list of libraries accessible by this user is required.

C A complete list of all libraries is requested.

AUTHSEL (field 9)

This field contains a value that selects only those libraries with a given access authority. The format is a coded value.

This value: Indicates:

R Only libraries for which you have read access are to be selected.

W Only libraries for which you have update access are to be selected.

FORWLOCT (field 10)

This field contains a command chaining value, which is used to receive the next response from Information Exchange (from the last breakpoint). To use this field:

1. Issue the initial List Libraries command with all blanks in this field.
2. Issue subsequent List Libraries commands with FORWLOCT set to NEXTLOC (field 8) of the previous List Libraries response, until you receive a response that indicates the end of the data. You reach the end of the data when CONTINUE (field 7) of a previous List Libraries response is 0.

BACKLOCT (field 11)

This field contains a command chaining value, which is used to receive the next response from Information Exchange (from the last breakpoint). To use this field, issue the List Libraries command with BACKLOCT set to NEXTLOC (field 8) of the previous List Libraries response, until you encounter the end of the list when CONTINUE (field 7) of a previous List Libraries response is 0.

DIRINDC (field 12)

This field contains a value that indicates the direction of the list.

This value: Indicates:

F List libraries in a forward direction from the forward locator FORWLOCT.

B List libraries in a backward direction from the backward locator BACKLOCT.

List Libraries response

This is Information Exchange's response to the List Libraries command; its length is 56 bytes, plus the length of the library member item.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDILLST
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Command return code
7	36	1	CONTINUE	List continued indicator
8	37	16	NEXTLOC	Locator value
9	53	4	ITEMS	Number of items
10	57	n	item one	Beginning of first library item

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator from the List Libraries command.

RETCODE (field 6)

This field contains error code values for expansion level 1. See "Handling library command responses" on page 189.

This field contains error code values for expansion level 2.

This value:	Indicates:								
06	Information Exchange did not find any libraries that fit the search criteria. No messages are returned to the mailbox.								
08	The session started with a message size that is too small. Check the mailbox for an error message.								
09	Information Exchange encountered an invalid command parameter. Check the mailbox for an error message for any of the following:								
	<table> <thead> <tr> <th>This Field:</th> <th>Indicates:</th> </tr> </thead> <tbody> <tr> <td>SELECT (field 8)</td> <td>You made an invalid selection. Select an <i>A</i> or <i>C</i>. Check the mailbox for an error message.</td> </tr> <tr> <td>AUTHSEL (field 9)</td> <td>You specified an invalid value. Select an <i>R</i> or a <i>W</i>. Check the mailbox for an error message.</td> </tr> <tr> <td>DIRINDC (field 12)</td> <td>You specified an invalid value. Select an <i>F</i> or a <i>B</i>. Check the mailbox for an error message.</td> </tr> </tbody> </table>	This Field:	Indicates:	SELECT (field 8)	You made an invalid selection. Select an <i>A</i> or <i>C</i> . Check the mailbox for an error message.	AUTHSEL (field 9)	You specified an invalid value. Select an <i>R</i> or a <i>W</i> . Check the mailbox for an error message.	DIRINDC (field 12)	You specified an invalid value. Select an <i>F</i> or a <i>B</i> . Check the mailbox for an error message.
This Field:	Indicates:								
SELECT (field 8)	You made an invalid selection. Select an <i>A</i> or <i>C</i> . Check the mailbox for an error message.								
AUTHSEL (field 9)	You specified an invalid value. Select an <i>R</i> or a <i>W</i> . Check the mailbox for an error message.								
DIRINDC (field 12)	You specified an invalid value. Select an <i>F</i> or a <i>B</i> . Check the mailbox for an error message.								
16	Information Exchange could not find the owning account specified in LIBACCT (field 7). No messages are returned to the mailbox.								

CONTINUE (field 7)

This field contains a value that indicates whether the library list continues.

This value:	Indicates:
0	The system has encountered the end of the data, and no more libraries are to follow.
1	More libraries are to follow.

NEXTLOC (field 8)

This field contains the value the system uses to submit a subsequent List Libraries command that continues with the next library entry in the directory by copying it into FORWLOCT (field 10) or BACKLOCT (field 11).

ITEMS (field 9)

This field contains the count of library entries that are to follow.

item one (field 10)

This field contains the library statistics and name. See “Library Entry item structure” on page 207 for a description of the Library entry item structure.

Library Entry item structure

The following is the format of the Library Entry item; its length is 173 bytes. This Library Entry item is returned as part of the List Libraries response. See “item one (field 10)” on page 206.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	LIBNAME	Library name
2	9	8	LIBACCT	Library owner account ID
3	17	8	LIBUSER	Library owner user ID
4	25	6	CDATE	Library creation date
5	31	6	CTIME	Library creation time
6	37	8	ULIBACCT	Library update account
7	45	8	ULIBUSER	Library update user ID
8	53	6	UUPDATE	Library update date
9	59	6	UTIME	Library update time
10	65	10	MEMCOUNT	Number of members
11	75	1	WRITETYP	Write authority level
12	76	8	WRITELST	Write list
13	84	1	READTYP	Read authority level
14	85	8	READLIST	Read list
15	93	1	SEARCH	Searchable indicator
16	94	1	OWNRESP	Owner responsible for retrieve
17	95	79	LIBTITLE	Library descriptive title

LIBNAME (field 1)

This field contains the name of the library for which the statistics and information pertain. The format is alphanumeric, left-justified, and padded on the right with blanks.

LIBACCT (field 2)

This field contains the account ID of the library owner. The format is alphanumeric, left-justified, and padded on the right with blanks.

LIBUSER (field 3)

This field contains the user ID of the library owner. The format is alphanumeric, left-justified, and padded on the right with blanks.

CDATE (field 4)

This field contains the date the library was defined, offset to the time zone specified in the Session Start command. The format is *YYMMDD*.

CTIME (field 5)

This field contains the time the library was defined, offset to the time zone specified in the Session Start command.

ULIBACCT (field 6)

This field contains the account ID of the user who last changed library information.

ULIBUSER (field 7)

This field contains the user ID of the user who last changed library information.

UPDATE (field 8)

This field contains the date the library information was changed, offset to the time zone specified in the Session Start command. The format is *YYMMDD*.

UTIME (field 9)

This field contains the time the library information was changed, offset to the time zone specified in the Session Start command. The format is *HHMMSS*.

MEMCOUNT (field 10)

This field contains the number of members that currently define and comprise this library.

WRITETYP (field 11)

This field contains the authority type for update access to the library.

This value:	Indicates:
P	Only the owner can update this library.
O	Only users within the same account can update this library.
G	Any user can update this library.
L	Any user in the list, which the WRITELST (field 12) specification defines, can update this library.

WRITELST (field 12)

This field contains the name of a permanent distribution list that details the users who can update the library.

READTYP (field 13)

This field contains the authority type for read access to the library. The format is a coded value, which can be any one of the following:

This value	Indicates
P	Only the owner can read this library.
O	Only users within the same account can read this library.
G	Any user can read this library.
L	Any user in the list, which READLIST (field 14) defines, can read this library.

READLIST (field 14)

This field contains the name of a permanent distribution list that details the users who have read access to the library.

SEARCH (field 15)

A *Y* in this field specifies that this library has been defined as a searchable library. An *N* specifies the library has been defined as nonsearchable.

OWNRESP (field 16)

This field contains a value that indicates whether the owner account is responsible for the charges associated with Library Retrieve commands.

This value: Indicates:

Y The receive-side charges associated with subsequent retrieve commands are charged to the library owner.

N The receive-side charges associated with subsequent retrieve commands are charged to the account that issued the Library Retrieve command.

LIBTITLE (field 17)

This field contains a description of the library.



NOTE: If a user who does not have access to a library requests library information, the system returns only the LIBNAME, LIBACCT, LIBUSER, and LIBTITLE fields. Information Exchange returns the remaining fields blank.

List Library Members command—SDILMBR

This command provides a list of library members or member statistics. The length of this command is 73 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILMBR
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Return response name
7	42	8	OWNER	Owning account ID
8	50	8	LIBNAME	Library name
9	58	16	LOCATOR	Locator value

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1 or 2. Invalid values default to 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the List Library Members response. The format is alphanumeric, left-justified, and padded on the right with blanks.

OWNER (field 7)

This field identifies the owner of the library. The default value is the account to which you belong. The format is alphanumeric, left-justified, and padded on the right with blanks.

LIBNAME (field 8)

This field specifies the name of the library from which you want to list the members. The format is alphanumeric, left-justified, and padded on the right with blanks.

LOCATOR (field 9)

This field contains a command chaining value that the system uses to receive the next response from Information Exchange, from the last breakpoint. To use this field:

1. Issue a List Library Members command with all blanks in this field.
2. Issue a List Library Members command with LOCATOR set to NEXTLOC (field 8) of the previous List Library Members response until the CONTINUE indicator is 0.

List Library Members response

This is Information Exchange's response to the List Library Members command; its length is 56 bytes, plus the length of the library member.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDILLST
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Command return code
7	36	1	CONTINUE	List continued indicator
8	37	16	NEXTLOC	Locator value
9	53	4	ITEMS	Number of items
10	57	n	item one	Beginning of first library item

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1. Invalid values default to 1.

RETCODE (field 6)

This field contains error code values for expansion level 1. See "Handling library command responses" on page 189.

This field contains error code values for expansion level 2.

This value: Indicates:

- | | |
|----|---|
| 06 | Information Exchange found no members for the library you specified. No messages are returned to the mailbox. |
| 10 | You do not have read-access authority for the library you specified in LIBNAME (field 8). No messages are returned to the mailbox. |
| 16 | You specified an invalid or unknown account ID in OWNER (field 7). No messages are returned to the mailbox. |
| 18 | You specified an invalid or unknown library name in LIBNAME (field 8). No messages are returned to the mailbox. |
| 30 | The read-access list for the library name specified in LIBNAME (field 8) does not exist. Contact your service administrator. No messages are returned to the mailbox. |

CONTINUE (field 7)

This field contains a value that indicates whether the member list continues.

This value: Indicates:

0 The system has encountered the end of the data, and no more members are to follow.

1 More members are to follow.

NEXTLOC (field 8)

This field contains the value the system uses to submit a subsequent List Library Members command that continues with the next member entry in the directory by copying it into LOCATOR (field 9).

ITEMS (field 9)

This field contains the number of library member list entries that are to follow.

item one (field 10)

This field contains the library member statistics and name. See “Library Member List entry” on page 213 for the structure of a member entry.

Library Member List entry

The format of the Library Member List entry is shown below; its length is 159 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	MEMNAME	Member name
2	9	20	CREATOR	User ID of the creator
3	29	6	CREDATE	Date created
4	35	6	CRETIME	Time created
5	41	20	UPDATER	User ID of the last to update
6	61	6	UPDDATE	Date last updated
7	67	6	UPDTIME	Time last updated
8	73	8	MEMLNTH	Member length
9	81	79	MEMDESC	Member description

MEMNAME (field 1)

This field contains the name of the member. The format is alphanumeric, left-justified, and padded on the right with blanks.

CREATOR (field 2)

This field contains the fully qualified user ID of the library member creator. The format is alphanumeric, left-justified, and padded on the right with blanks.

CREDATE (field 3)

This field contains the date the member was created, offset to the time zone specified in the Session Start command. The format is *YYMMDD*.

CRETIME (field 4)

This field contains the time the member was created, offset to the time zone specified in the Session Start command. The format is *HHMMSS*.

UPDATER (field 5)

This field contains the fully qualified user ID of the last user to update this member.

UPDDATE (field 6)

This field contains the date the member was last updated, offset to the time zone specified in the Session Start command. The format is *YYMMDD*.

UPDTIME (field 7)

This field contains the time the member was last updated, offset to the time zone specified in the Session Start command. The format is *HHMMSS*.

MEMLNTH (field 8)

This field contains the length of the member in bytes.

MEMDESC (field 9)

This field contains the title or description associated with this member.

Retrieve Library Member command—SDILRTV

This command places a copy of a library member in your mailbox. Information Exchange places the member in the mailbox as a message group. Library members can be directed to user mailboxes using the standard Information Exchange destination specification. Charges are based on receive domain payment levels. The length of this command is 112 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILRTV
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Return response name
7	42	8	OWNER	Owning account ID
8	50	8	LIBNAME	Library name
9	58	8	MEMBER	Member name
10	66	2	MSGRETN	Message retention
11	68	8	MSGNAME	Message name
12	76	5	MSGSEQN	Message sequence
13	81	8	MSGUCLS	Message user class
14	89	1	MSGCLASS	Message delivery class (blank, <i>P</i> , <i>I</i>)
15	90	1	MSGCHRG	Message service charge (<i>1</i> , <i>3</i> , <i>5</i> , <i>6</i>)
16	91	1	MSGRCPTS	Message acknowledgment (blank, <i>R</i> , <i>D</i> , <i>B</i> , <i>A</i> , <i>C</i> , <i>E</i> , <i>F</i>)
17	92	8	DESTACCT	Destination account or distribution list or first 8 characters of destination alias name
18	100	8	DESTUSER	Destination user ID or last 8 characters of destination alias name
19	108	1	DESTTYPE	Destination type (<i>D</i> , <i>L</i>)
20	109	1	DTBLTYPE	Alias table type (<i>G</i> , <i>O</i> , <i>P</i>) or intersystem indicator (<i>I</i>)
21	110	3	DTBLID	Alias table name or intersystem system ID

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1 or 2. Invalid values default to 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Retrieve Library Member response. The format is alphanumeric, left-justified, and padded on the right with blanks.

OWNER (field 7)

This field contains a value that identifies the owner of the library specified by LIBNAME (field 8). The default value is the account to which you belong. The format is alphanumeric, left-justified, and padded on the right with blanks.

LIBNAME (field 8)

This field contains the name of the library where you want to retrieve the specified member. The format is alphanumeric, left-justified, and padded on the right with blanks.

MEMBER (field 9)

This field contains the name of the member you want to retrieve. The format is alphanumeric, left-justified, and padded on the right with blanks.

MSGRETN (field 10)

This field contains a value that specifies the number of days that Information Exchange keeps the message group stored if it is not received by the destination. If this field does not contain valid numeric digits, or if 00 is specified, the system supplies a default value. The format is numeric digits. The maximum value that may be specified is 180 days. Information Exchange ignores any value greater than the maximum and uses the default.



NOTE: The field size will remain at 2 digits and the user is responsible for encoding the number of days as follows:

The high-order digit may be 0-9 or A-I. If A-I is used, it is translated to mean 10-18. The low-order digit may be 0-9 and will remain translated as 0-9. That is, the user would supply a value of *F6* to indicate a message retention of 156 days.

MSGNAME (field 11)

This field contains a name that is given to a message or group of messages. This is applied to the member data the system stores.

MSGSEQN (field 12)

This field contains a message input sequence number.

MSGUCLS (field 13)

This field contains a user-defined classification for the message. It can be referenced with the Receive Message command to select which messages on your queue you want to receive.

MSGCLASS (field 14)

This field contains a value that indicates the class of delivery for the message.

This value:	Indicates:
blank	A normal-priority delivery.
P	A high-priority delivery.
I	An express-priority delivery.

MSGCHRG (field 15)

This field contains a value that indicates to Information Exchange how the retriever wants the message charges to be paid.

Message charge code:	Means:
1	The receiver pays all the charges for the messages.
3	The receiver pays the charges. If the receiver does not agree to pay the charges, the sender pays all the charges.
5	The sender pays the charges. If the sender does not pay the charges, the receiver pays the charges.
6	The sender pays the charges.

Figure 8. Values of MSGCHRG and their meaning

For more information, refer to *Information Exchange Charges Reference*.

MSGRCPTS (field 16)

This field contains a value that indicates the type of acknowledgment messages you want to receive. The value placed here is the same as the value in the Send Message command long header. See “Send Message command—SDISNDM” on page 134.

DESTACCT (field 17)

The format is alphanumeric, left-justified, and padded on the right with blanks. See “Understanding destination address specification” on page 141.

DESTUID (field 18)

The format is alphanumeric, left-justified, and padded on the right with blanks. See “Understanding destination address specification” on page 141.

DESTTYPE (field 19)

This field contains an address specification:

This value:	Indicates:
D	True local ID, intersystem ID, or alias reference.
L	The name of a distribution list.

See “Understanding destination address specification” on page 141.

DTBLTYP (field 20)

See “Understanding destination address specification” on page 141.

DTBLID (field 21)

This field is alphanumeric, left-justified, and padded on the right with blanks. For more information on alias tables, See “Understanding destination address specification” on page 141.

Retrieve Library Member response

This is Information Exchange's response to the Retrieve Library Member command; its length is 35 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDILRTV
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Command return code

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator from the Retrieve Library Member command. See "Handling library command responses" on page 189.

RETCODE (field 6)

This field contains error code values for expansion level 1. See "Handling library command responses" on page 189.



NOTE: Errors, warnings, and acknowledgments are placed in the library owner's mailbox, if the owner is willing to pay for retrieval. Otherwise, these messages are placed in the requester's mailbox. Issue a Receive Message command to retrieve any messages.

This field contains error code values for expansion level 2.

This value: Indicates:

- | | | | | | | | | | | | |
|---------------------|---|--------|------------|---------------------|---|--------------------|--|---------------------|---|---------------------|--|
| 09 | An invalid value was specified for one or more of the following fields. Check the mailbox for one or more warning messages. | | | | | | | | | | |
| | <table border="0" style="width: 100%;"> <tr> <td style="width: 30%; padding-right: 20px;">Field:</td> <td>Indicates:</td> </tr> <tr> <td>MSGCLASS (field 14)</td> <td>You specified an invalid or unknown message delivery class.</td> </tr> <tr> <td>MSGCHRG (field 15)</td> <td>This field contained an invalid message service charge value. See “MSGCHRG (field 17)” on page 139 for more information.</td> </tr> <tr> <td>MSGRCPTS (field 16)</td> <td>This field contained an invalid value and does not match the value in the MSGRCPTS (field 17) of the Send Message command. See “Send Message command—SDISNDM” on page 134 for more information.</td> </tr> <tr> <td>DESTTYPE (field 19)</td> <td>You specified an invalid destination type.</td> </tr> </table> | Field: | Indicates: | MSGCLASS (field 14) | You specified an invalid or unknown message delivery class. | MSGCHRG (field 15) | This field contained an invalid message service charge value. See “MSGCHRG (field 17)” on page 139 for more information. | MSGRCPTS (field 16) | This field contained an invalid value and does not match the value in the MSGRCPTS (field 17) of the Send Message command. See “Send Message command—SDISNDM” on page 134 for more information. | DESTTYPE (field 19) | You specified an invalid destination type. |
| Field: | Indicates: | | | | | | | | | | |
| MSGCLASS (field 14) | You specified an invalid or unknown message delivery class. | | | | | | | | | | |
| MSGCHRG (field 15) | This field contained an invalid message service charge value. See “MSGCHRG (field 17)” on page 139 for more information. | | | | | | | | | | |
| MSGRCPTS (field 16) | This field contained an invalid value and does not match the value in the MSGRCPTS (field 17) of the Send Message command. See “Send Message command—SDISNDM” on page 134 for more information. | | | | | | | | | | |
| DESTTYPE (field 19) | You specified an invalid destination type. | | | | | | | | | | |
| 10 | You do not have read-access authority for the library name you specified in LIBNAME (field 8). No messages are returned to the mailbox. | | | | | | | | | | |
| 15 | You specified an unknown or invalid system ID in DTBLID (field 21). No messages are returned to the mailbox. | | | | | | | | | | |
| 16 | You specified an unknown or invalid account ID in DESTACCT (field 17). No messages are returned to the mailbox. | | | | | | | | | | |
| 17 | You specified an unknown or invalid destination user ID in DESTACTT (field 17). No messages are returned to the mailbox. | | | | | | | | | | |
| 18 | You specified an unknown or invalid library name in LIBNAME (field 8). No messages are returned to the mailbox. | | | | | | | | | | |
| 19 | You specified an unknown or invalid member name in MEMBER (field 9). No messages are returned to the mailbox. | | | | | | | | | | |
| 23 | You are not authorized to retrieve library members to the user ID specified. No messages are returned to the mailbox. | | | | | | | | | | |
| 30 | The read-access list specified for this library does not exist. Contact your service administrator. No messages are returned to the mailbox. | | | | | | | | | | |
| 31 | The retrieved destination list referenced in DESTACCT (field 17) is missing. No messages are returned to the mailbox. | | | | | | | | | | |
| 32 | You specified an alias name in DTBLTYPE (field 20) and DTBLID (field 21) which is invalid or unknown. No messages are returned to the mailbox. | | | | | | | | | | |
| 33 | You are not able to retrieve messages because of an ESO conflict. No messages are returned to the mailbox. | | | | | | | | | | |

Retrieve Library Member CDH command—SDIRCDH

This command requests the CDH information for a library member. The length of this command is 70 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIRCDH
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	8	LIBACCT	Library owning account
8	50	8	LIBRARY	Library name
9	58	8	MEMBER	Member name
10	66	5	LOCATOR	CDH offset

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the single character 1 or 2. The default value for the expansion-level indication is 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Retrieve Library Member CDH response.

LIBACCT (field 7)

This field contains the account ID of the account that owns the library being referenced.

LIBRARY (field 8)

This field contains the library name being referenced.

MEMBER (field 9)

This field contains the library member name being referenced.

LOCATOR (field 10)

This field is used to retrieve CDHs that are too long to fit within your maximum transmission size. Use this field to indicate that you want to browse the remaining records. Place an offset (relative to zero) in this field to retrieve CDH data starting from that position. The field should be right-justified and padded on the left with zeros. If left blank, a value of zeros is assumed.

Submit the initial Retrieve Library Member CDH command with blanks in this field. Submit subsequent Retrieve Library Member CDH commands with the NEXTLOC value from the previous Retrieve Library Member CDH response in this field, until you receive a response that contains “00000” in the NEXTLOC field.

Retrieve Library Member CDH response

This is Information Exchange's response to the Retrieve Library Member CDH command. The length of this response is variable (based on the length of the portion of the CDH returned).

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIRCDH
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RESPONSE	Response code
7	36	1	CONTINUE	Continuation flag
8	37	5	NEXTLOC	Offset to next CDH position
9	42	5	LENGTH	Length
10	47	n	CDH	CDH data

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the value from EXPAND (field 5) of the Retrieve Library Member CDH command.

RESPONSE (field 6)

This field contains error code values. See "Browse Library Member command—SDILBRW" on page 190.

If the expansion level is 2, this field contains one of the following values:

This value:	Indicates:
00	Normal execution was completed.
10	Read access to the library is not allowed.
11	The offset field contains other than numeric characters or blanks.
12	The offset field is greater than the length of the CDH.
16	The owning account was not found.
18	The library was not found.
19	The library member was not found.
30	The list of users allowed to access the library was not found.

CONTINUE (field 7)

This field contains one of the following:

This value: Indicates:

0 The last byte of the CDH data was transmitted.

1 There is more CDH data to transmit.

NEXTLOC (field 8)

The NEXTLOC field contains the beginning offset to the remaining CDH data. Use this field to submit a subsequent Retrieve Library Member CDH command that continues with the next byte of the CDH data by copying the value into the LOCATOR field. If the CONTINUE flag is set to 0, this field contains "00000."

LENGTH (field 9)

This field contains the length of the following CDH (field 10) data.

CDH (field 10)

This field contains the CDH data for the library member.

Search Library command—SDILSCH

This command returns a list of library members that meet the requirements of the search string. The length of this command is 62 bytes, plus the length of the search argument string.

Use the Library Title command following a Search Library command to retrieve the additional members found during the search that could not be transmitted in the single buffer allocated for the search response.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDILSCH
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Return response name
7	42	8	SRCHACCT	Account to search
8	50	8	SRCHLIB	Library to search
9	58	1	SORTSEQ	Sort sequence (<i>A, D</i>)
10	59	4	LENGTH	Length of search argument string
11	63	var	SRCHARG	Search argument string

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1 or 2. Invalid values default to 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Search Library response. The format is alphanumeric, left-justified, and padded on the right with blanks.

SRCHACCT (field 7)

This field contains the owning account of the library you want to search.

SRCHLIB (field 8)

This field contains the name of the library you want to search.

SORTSEQ (field 9)

This field indicates the order in which the members are to be returned.

This value: Indicates:

A Ascending creation date and time.

D Descending creation date and time.

Any other value indicates that no date and time sort is to be performed.

LENGTH (field 10)

This field contains the length of the search arguments.

SRCHARG (field 11)

This field contains the search argument string whose length is specified in LENGTH (field 10).

A search argument string is a character string that represents the words to be located and logical relationships to be used (such as or, and, or not). The search string expression is in an algebraic format. A word can contain alphanumeric characters, plus the following special characters: = (X'7E'), # (X'7B'), \$(X'5B'), % (X'6C'), or @ (X'7C').

During the scan for words, the system discards other characters in a word. A word is delimited with a blank or one of the following characters, which are the logical operators:

This value:	Indicates:
X'4D' or X'AD' (left parenthesis or left bracket)	The beginning of a logical expression.
X'5D' or X'BD' (right parenthesis or right bracket)	The end of a logical expression.
X'60' (dash or minus sign)	A range of words.
X'5C' (asterisk)	Word truncation or abbreviation.
X'50' (ampersand)	The logical and operator.
X'5F' or X'A1' (not sign or tilde)	The logical not operator.
X'4F', X'5A' or X'6A' (vertical bar, exclamation point, or split vertical bar)	The logical or operator.

For example:

(CAT&~DOG) HOR SE*	This means the system finds documents with words beginning with horse or documents containing the word cat, and not containing the word dog.
X*-Z*	This means the system finds all documents with words which begin with an X, Y, or Z.

Search Library response

Information Exchange returns the following response to the Search Library command; its length is 45 bytes, plus the length of the library item member.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDILSCH
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Command return code
7	36	5	NEXTLOC	Locator value
8	41	5	ITEMS	Number of items
9	46	n	item one	Beginning of first member item

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator from the search command.

RETCODE (field 6)

This field contains error code values for expansion level 1. See “Handling library command responses” on page 189.

This field contains error code values for expansion level 2.

This value:	Indicates:
08	Information Exchange encountered one or more internal errors. Contact Customer Care. Check the mailbox for error messages.
09	You specified a nonnumeric length in LENGTH (field 10). Check the mailbox for messages.
10	You do not have access authority to the library you specified in SRCHLIB (field 8). No messages are returned to the mailbox.
16	You specified an unknown or invalid account ID in SRCHACCT (field 7). No messages are returned to the mailbox.
18	You specified an unknown or invalid library name in SRCHLIB (field 8). No messages are returned to the mailbox.
30	The read authorization list specified for this library does not exist. Contact your service administrator. No messages are returned to the mailbox.
40	You need to enter more specific criteria in SRCHARG (field 11). More than 1000 members were found. No messages are returned to the mailbox.
41	Your search criteria contained an invalid combination of search arguments in SRCHARG (field 11). No messages are returned to the mailbox.
42	Your search string contained invalid data after the operand in SRCHARG (field 11). No messages are returned to the mailbox.
43	Your search string contained invalid data after the range operator in SRCHARG.
44	Your search string contained invalid data after the non-range operator in SRCHARG (field 11). No messages are returned to the mailbox.
45	The parentheses in your search criteria in SRCHARG (field 11) are not balanced. No messages are returned to the mailbox.
46	Your first search argument is alphabetically greater than your second search argument in SRCHARG (field 11). No messages are returned to the mailbox.

NEXTLOC (field 7)

This field is used to submit a subsequent Library Title command that continues with the next library entry in the directory by copying it into LOCATOR (field 7). If this field contains a series of Fs, there are no more entries.

ITEMS (field 8)

This field contains the number of member entries that are to follow.

ITEM ONE (field 9)

This field contains the library member statistics and name that meet your search criteria. See “Library Member item format” on page 230 for the structure of a member item.

Library Title command—SDITITL

This command returns a list of library members from a previous search command; its length is 49 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDITITL
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Return response name
7	42	5	LOCATOR	Locator value
8	47	3	COUNT	Number of titles to return

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1 or 2. Invalid values default to 1. See “Handling library command responses” on page 189.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Library Title response. The format is alphanumeric, left-justified, and padded on the right with blanks.

LOCATOR (field 7)

This field contains a command chaining value, which the system uses to receive the next response from the Search Library command (from the last breakpoint). To use this field:

- Issue a Library Title command with LOCATOR set to NEXTLOC (field 7) of the previous Search Library response.
- Issue a Library Title command with LOCATOR set to the NEXTLOC (field 7) of the previous Library Title response, until the field contains a series of Fs.

COUNT (field 8)

This field contains the maximum number of library member items to return.

Library Title response

This is Information Exchange's response to the Library Title command; its length is 45 bytes, plus the length of the library member item.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDITITL
2	9	8	ACCOUNT	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Command return code
7	36	5	NEXTLOC	Locator value
8	41	5	ITEMS	Number of items
9	46	n	item one	Beginning of first member item

COMMAND (field 1)

ACCOUNT (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the command expansion-level indicator. It must contain the value 1. Invalid values default to 1.

RETCODE (field 6)

This field contains error code values for expansion level 1. See "Handling library command responses" on page 189.

This field contains error code values for expansion level 2.

This value:	Indicates:
08	Information Exchange could not find the results from a previous search. Check the mailbox for an error message.
09	Information Exchange encountered an invalid command parameter in COUNT (field 8). Check the mailbox for an error message.
11	Information Exchange either encountered an invalid value in LOCATOR (field 7) or could not find the results from a previous search. No messages are returned to the mailbox.

NEXTLOC (field 7)

The system uses this field to submit a subsequent Library Title command, which continues with the next library entry in the directory by copying it into LOCATOR (field 7) of the Library Title command. If this field contains a series of Fs, there are no more entries.

ITEMS (field 8)

This field contains the count of member entries that follow.

item one (field 9)

This field contains the library member's statistics and names. See "Library Member item format" on page 230 for the structure of a member item.

Library Member item format

The format of the member item the system returns in a Library Search or Library Title response is shown below; its length is 100 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	LIBRARY	Library name
2	9	8	MEMBER	Member name
3	17	79	MEMDESC	Member description
4	96	5	FILLER	Filler

LIBRARY (field 1)

This field contains the name of the library.

MEMBER (field 2)

This field contains the name of the member.

MEMDESC (field 3)

This field contains the title or description that is associated with this member.

FILLER (field 4)

This field is not used and is set to blanks.

Audit Commands

Audit Browse command-SDIAUDB

With this command, you can browse the contents of your audit data. Information Exchange's response to this command is the Audit Browse response. This response contains the audit records that provide information on the status of messages. The length of the Audit Browse command is 141 bytes for expansion levels 1 and 2, 149 bytes for expansion levels 3 and 4.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIAUDB
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	1	RECTYPES	Type of records to retrieve (<i>A, C, S, R, B</i>)
8	43	8	ALTUSRID	Alternate user ID
9	51	20	TRADPART	Trading partner
10	71	7	DATEFROM	From date (CYMMDD)
11	78	7	DATETO	To date (CYMMDD)
12	85	8	MSGUCLS	User message class
13	93	1	STATUS	Status selection (<i>D, P U</i> , blank)
14	94	1	TIMEZONE	Time zone for dates
15	95	47	LOCATOR	Command chaining value
16	142	8	ALTACCT	Alternate account ID

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a command expansion-level indicator. Invalid values default to *blank* or 1.

This value:	Indicates:
blank or 1	Do not include extended timing fields in the audit record.
2	Include extended timing fields in the audit record.
3	Include extended timing fields and EDI exchange control number in the audit record. When expansion level 3 is specified, you can request audits for a user or users in an account ID other than your own.

This value:	Indicates:
4	Include extended timing fields, EDI exchange control number fields, carbon copy fields, and other miscellaneous fields. When expansion level 4 is specified, you can request audits for a user or users in an account ID other than your own. To request carbon copy as partner records, expansion level 4 must be specified.

RETRSPC (field 6)

This field contains the value that is to be returned in the first 8 characters of the Audit Retrieve response.

RECTYPES (field 7)

This field specifies the type of message records that Information Exchange returns. The format is a coded value.

This value:	For expansion levels:	Indicates:
S	1, 2, 3, 4	Messages that have been sent by the user.
R	1, 2, 3, 4	Messages that are receivable have been received or have been purged before receipt (see status, field 13).
B	1, 2, 3, 4	Messages that have been both sent and received by the user. This is the default.
A	4 only	All audit types (sent, received/receivable, and carbon copy as partner).
C	4 only	Carbon copy as partner.

ALTUSRID (field 8)

This field, together with ALTACCT (field 16), allows you to request audit records for an account and user ID other than your own or for a cluster parent user ID. You must be authorized to request audit records other than your own. Information Exchange ignores input from non authorized users. If you leave this field blank or you are not authorized to view audit records for the alternate account and user ID, Information Exchange extracts only your own audit records. If a value is specified in this field but the ALTACCT (field 16) is left blank, Information Exchange will extract audit records for the account specified in the ACCNTNO (field 2) and the user ID specified in this field. To request audit records for all users in the specified account, enter a question mark (?) in this field.

TRADPART (field 9)

This field should contain the ID of the trading partner for whom you want audit records of messages. If you do not want to select records for a specific trading partner, leave this field blank. A trading partner can be specified directly by using the trading partner's Information Exchange ID, or indirectly by using an already defined alias reference for the trading partner.

- For send audits, this field represents the receiver of the message.
- For receive audits, this field represents the sender of the message.
- For carbon copy as partner audits, this field represents the requester of the carbon copy message.

ISSSAAAAAAAAUUUUUUUU indicates a trading partner's Information Exchange ID where:

This value:	Indicates:
I	This is a remote trading partner. For local trading partners, leave blank.
SSS	The remote system ID for a remote trading partner. For local trading partners, leave blank
AAAAAAA	The account ID of the trading partner.
UUUUUUUU	The user ID of the trading partner. To request audit records for all trading partners in an account, enter a question mark(?) in this position of the trading partner field.

TNNNAAAAAAAAAAAAAAAA indicates a trading partner's alias reference where:

This value:	Indicates:
T	The alias table type. Values <i>G</i> , <i>P</i> , and <i>O</i> are allowed in this field.
NNN	The alias table name.
AAAAAAAAAAAAAA AA	The alias name of the trading partner.

DATEFROM (field 10)

Information Exchange extracts audit records representing messages sent or received on or after the date you enter in this field, offset to the value specified in TIMEZONE (field 14). Enter the date in the format *CYYMMDD*, where *C* is the century (0=1900, 1=2000), *YY* is the year, *MM* is the month, and *DD* is the day of the month. For example, 0890502 represents May 2, 1989. The default is 0000102 (January 2, 1900).

For sent audits and carbon copy as partner audits:

Information Exchange extracts audit records representing messages sent on or after the date you enter in this field. This is based on the value in the date sent field ASNDDATE (field 20) of the Audit Record/Display format.

For receivable audits:

If the message has been received, Information Exchange extracts audit records representing messages received on or after the date you enter in this field. This is based on the value in the date purged/delivered field ARCVDATE (field 22) of the Audit Record/Display format.

If the message has not been received, Information Exchange extracts audit records representing messages sent on or after the date you enter in this field. This is based on the value in the date sent field ASNDDATE (field 20) of the Audit Record/Display format.

DATETO (field 11)

Information Exchange extracts audit records representing messages sent or received on or before the date you enter in this field, offset to the value specified in TIMEZONE (field 14). Information Exchange extracts audit records representing messages sent or received on or before the date you enter in this field. Enter the date in the format *CYYMMDD*, where *C* is the century (0=1900, 1=2000), *YY* is the year, *MM* is the month, and *DD* is the day of the month. For example, 1010502 represents May 2, 2001. The default is 1420916 (September 16, 2042).

For sent audits and carbon copy as partner audits:

Information Exchange extracts audit records representing messages sent on or before the date you enter in this field. This is based on the value in the date sent field ASNDDATE (field 20) of the Audit Record/Display format.

For receivable audits:

If the message has been received, Information Exchange extracts audit records representing messages received on or before the date you enter in this field. This is based on the value in the date purged/delivered field ARCVDATE (field 22) of the Audit Record/Display format.

If the message has not been received, Information Exchange extracts audit records representing messages sent on or before the date you enter in this field. This is based on the value in the date sent field ASNDDATE (field 20) of the Audit Record/Display format.

MSGUCLS (field 12)

This field contains the user-message class of the messages you want to select. You can use a question mark (?) to substitute for any character in any position. The field is left-justified and padded on the right with blanks.

STATUS (field 13)

This field contains a value that specifies which messages should be selected. The format is a coded value.

This value: Indicates:

D A status of D can be one of the following:

- Received
- Sent to a non-Information Exchange system
- Transferred
- Redirected

P • Purged

U A status of U can be one of the following:

- In transit to another system
- Queued
- Forwarded
- Mailboxed

blank • All messages regardless of their status. This is the default.

TIMEZONE (field 14)

This field indicates the time zone for the dates in DATEFROM (field 10) and DATETO (field 11). Enter either *L* (for the time zone specified in the Session Start command) or *G* (for Greenwich mean time). The default is *L*.

LOCATOR (field 15)

This field contains a command chaining value that receives information for additional audit records. Because of the maximum message size specifications you indicated in MAXMSGSZ (field 10) of the Session Start command, Information Exchange may not have been able to send you all available audit records. Use this field to indicate that you want to continue the browse process.

Submit the initial Audit Browse command with blanks in this field. Submit subsequent Audit Browse commands with the CHAINRSP value from the previous Audit Browse response in this field, until you receive a response that contains a series of 9s in the CHAINRSP field. A series of 9s indicates all audit items were returned to you.

ALTACCT (field 16)

This field allows you to request audit records for an account and user ID other than your own. If left blank, audit records for the account specified in ACCNTNO (field 2) are returned. If a value is specified in this field, ALTUSRID (field 8) must also be specified. The requester must have authority to view audit records for the alternate account and user ID.



NOTE: This field is only valid for expansion levels 3 and 4 requests.

Audit Browse response

This is Information Exchange's response to the Audit Browse command. This response contains the audit records that provide information on the status of your messages. The length of this response is 85 bytes plus the amount of data in the audit items.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIAUDB
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	47	CHAINRSP	Command chaining value
7	81	2	ERRCODE	Error code
8	83	3	ITEMS	Number of audit items
9	86	n	AUDIT ITEMS	Beginning of first audit item

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field displays the value from EXPAND (field 6) of the Audit Browse command.

CHAINRSP (field 6)

This field is a command chaining value to be copied into the LOCATOR field of subsequent Audit Browse commands. If this value contains a series of 9s, there are no more audit records that meet your browse specifications.

ERRCODE (field 7)

A value of *00* in this field indicates that no errors occurred. A value of *04* indicates there is a recoverable error. A value of *08* indicates there is a nonrecoverable error present. If this field contains a value other than *00*, issue a Receive Message command to retrieve the associated error messages.

ITEMS (field 8)

This field contains the number of audit items that follow.

AUDIT ITEMS (field 9)

See "Audit Record/Display format" on page 243 for a description of the audit record display format.

Audit Retrieve command—SDIAUDR

This command downloads the contents of your audit file, in the form of a message group, into your Information Exchange mailbox. When downloading is complete, you must receive these messages with the Receive Message command. If no audit records are found that meet the selection criteria, Information Exchange places a blank (zero-length) message in your Information Exchange mailbox. The length of the Audit Retrieve command is 97 bytes for expansion levels 1 and 2, 105 bytes for expansion levels 3 and 4.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIAUDR
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	1	RECTYPES	Types of records to retrieve (<i>C, A, S, R, B</i>)
8	43	8	ALTUSRID	Alternate user ID
9	51	20	TRADPART	Trading partner
10	71	7	DATEFROM	From date (CYMMDD)
11	78	7	DATETO	To date (CYMMDD)
12	85	8	MSGUCLS	User message class
13	93	1	STATUS	Status selection (<i>D, P, U, blank</i>)
14	94	1	TIMEZONE	Time zone for dates
15	95	3	MAXMSGSZ	Maximum message length
16	98	8	ALTACCT	Alternate account ID

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains a command expansion-level indicator. Invalid values default to blank or 1.

This value:	Indicates:
blank or 1	Do not include extended time fields in the audit record.
2	Include extended timing fields in the audit record.
3	Include extended timing fields and EDI exchange control number in the audit record. When expansion level 3 is specified, you can request audits for a user or users in an account other than your own.
4	Include extended timing fields, EDI exchange control number fields, carbon copy fields, and other miscellaneous fields. When expansion level 4 is specified, you can request audits for a user or users in an account ID other than your own. To request carbon copy as partner records, expansion level 4 must be specified.

RETRSPC (field 6)

This field contains the value that returns in the first 8 characters of the audit retrieval response.

RECTYPES (field 7)

This field specifies the type of message records that Information Exchange returns. The format is a coded value.

This value:	For expansion levels:	Indicates:
S	1, 2, 3, 4	Messages that have been sent by the user.
R	1, 2, 3, 4	Messages that are receivable have been received or have been purged before receipt (see <i>STATUS</i> , field 13).
B	1, 2, 3, 4	Messages that have been both sent and received by the user. This is the default.
A	4 only	All audit types (sent, received/receivable, and carbon copy as partner).
C	4 only	Carbon copy as partner.

ALTUSRID (field 8)

This field, together with *ALTACCT* (field 16), allows you to request audit records for an account and user ID other than your own or for a cluster parent user ID. You must be authorized to request audit records other than your own. Information Exchange ignores input from non authorized users. If you leave this field blank or you are not authorized to view audit records for the alternate account and user ID, Information Exchange extracts only your own audit records. If a value is specified in this field but the *ALTACCT* (field 16) is left blank, Information Exchange will extract audit records for the account specified in the *ACCNTNO* (field 2) and the user ID specified in this field. To extract audit records for all users in the specified account, enter a question mark (?) in this field.

TRADPART (field 9)

This field contains the ID of the trading partner for whom you want audit records of messages. If you do not want to select records for a specific trading partner, leave this field blank. A trading partner can be specified directly by using the trading partner's Information Exchange ID, or indirectly by using an already defined alias reference for the trading partner.

- For send audits, this field represents the receiver of the message.
- For receive audits, this field represents the sender of the message.
- For carbon copy as partner audits, this field represents the requester of the carbon copy message.

ISSSAAAAAAAAUUUUUUUU indicates a trading partner's Information Exchange ID where:

This value:	Indicates:
I	This is a remote trading partner. For local trading partners, leave blank.
SSS	The remote system ID for a remote trading partner. For local trading partners, leave blank.
AAAAAAA	The account ID of the trading partner.
UUUUUUUU	The user ID of the trading partner. To extend audit records for all trading partners in an account, enter a question mark(?) in this position of the trading partner field.

TNNNAAAAAAAAAAAAAAAA indicates a trading partner's alias reference where:

This value:	Indicates:
T	The alias table type. Values <i>G</i> , <i>P</i> , and <i>O</i> are allowed in this field.
NNN	The alias table name.
AAAAAAAAAAAAAA AA	The alias name of the trading partner.

DATEFROM (field 10)

Information Exchange extracts audit records representing messages sent or received on or after the date you enter in this field, offset to the value specified in TIMEZONE (field 14). Enter the date in the format *CYYMMDD*, where *C* is the century (0=1900, 1=2000), *YY* is the year, *MM* is the month, and *DD* is the day of the month. For example, 0890502 represents May 2, 1989. The default is 0000102 (January 2, 1900).

For sent audits and carbon copy as partner audits:

Information Exchange extracts audit records representing messages sent on or after the date you enter in this field. This is based on the value in the date sent field ASNDDATE (field 20) of the Audit Record/Display format.

For receivable audits:

If the message has been received, Information Exchange extracts audit records representing messages received on or after the date you enter in this field. This is based on the value in the date purged/delivered ARCVDAT (field 22) of the Audit Record/Display format.

If the message has not been received, Information Exchange extracts audit records representing messages sent on or after the date you enter in this field. This is based on the value in the date sent field ASNDDATE (field 20) of the Audit Record/Display format.

DATETO (field 11)

Information Exchange extracts audit records representing messages sent or received on or before the date you enter in this field, offset to the value specified in TIMEZONE (field 14). Information Exchange extracts audit records representing messages sent or received on or before the date you enter in this field. Enter the date in the format *CYYMMDD*, where *C* is the century (0=1900, 1=2000), *YY* is the year, *MM* is the month, and *DD* is the day of the month. For example, 1010502 represents May 2, 2001. The default is 1420916 (September 16, 2042).

For sent audits and carbon copy as partner audits:

Information Exchange extracts audit records representing messages sent on or before the date you enter in this field. This is based on the value in the date sent field ASNDDATE (field 20) of the Audit Record/Display format.

For receivable audits:

If the message has been received, Information Exchange extracts audit records representing messages received on or before the date you enter in this field. This is based on the value in the date purged/delivered field ARCVDATE (field 22) of the Audit Record/Display format.

If the message has not been received, Information Exchange extracts audit records representing messages sent on or before the date you enter in this field. This is based on the value in the date sent field ASNDDATE (field 20) of the Audit Record/Display format.

MSGUCLS (field 12)

This field contains the user message class of the messages you want to select. You can use a question mark (?) as a substitute for any character in any position. The field is left-justified and padded on the right with blanks.

STATUS (field 13)

This field contains a value that specifies which messages should be selected. The format is a coded value.

This value: Indicates:

- | | |
|-------|---|
| D | A status of D can be one of the following: <ul style="list-style-type: none"> • Received • Sent to a non-Information Exchange system • Transferred • Redirected |
| P | <ul style="list-style-type: none"> • Purged |
| U | A status of U can be one of the following: <ul style="list-style-type: none"> • In transit to another system • Queued • Forwarded • Mailboxed |
| blank | <ul style="list-style-type: none"> • All messages regardless of their status. This is the default. |

TIMEZONE (field 14)

This field indicates the time zone for the dates in DATEFROM (field 10) and DATETO (field 11). Enter either *L* (for your local time zone specified in the Session Start command) or *G* (for Greenwich mean time). The default is *L*.

MAXMSGSZ (field 15)

This field contains a number that indicates the size of the largest message to be created in a message group. Checkpoint-level recovery uses this number to take commit points while receiving the message group that Information Exchange builds. If the field contains zero or any nonnumeric characters, Information Exchange creates a single message in the message group. The value in this field is multiplied by 1000 to obtain the individual message size.

ALTACCT (field 16)

This field allows you to request audit records for an account/user ID other than your own. If left blank, audit records for the account specified in ACCNTNO (field 2) are returned. If a value is specified in this field, ALTUSRID (field 8) must also be specified. The requester must have authority to retrieve audit records for the alternate account/user ID.



NOTE: This field is only valid for expansion level 3 and 4 requests.

Output format

The Audit Retrieve command builds a message group in your Information Exchange mailbox. The following are in the message group's identification fields:

This field:	Contains:
DESTACCT	*SYSTEM*
DESTUID	*AUDITS*
DESTTYPE	D
MSGNAME	blanks
MSGSEQN	blanks
MSGNCLS	blank
MSGUCLS	#SAUDIT
MSGRCPTS	blank
MSGCHRG	5
SYSTYPE	IBMIE
SYSLEVEL	Current version and release level of Information Exchange
DTBLTYP	blank
DTBLID	blanks

The content of a logical record in the message group is illustrated in “Audit Record/Display format” on page 243.

Audit Retrieve response

This is Information Exchange's response to the Audit Retrieve command; its length is 35 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIAUDR
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	2	RETCODE	Response code

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the value from EXPAND (field 5) of the Audit Retrieve command.

RETCODE (field 6)

This field contains error code values. If the expansion level is 1, this field contains one of the following values:

This value:	Indicates:
00	Normal execution was completed.
04	A warning message was placed in your mailbox and the command was executed.
08	An error message was placed in your mailbox and the command was not executed.

Audit Record/Display format

The following record layout represents the format of the audit records as they exist within the Information Exchange message or the Audit Browse response. All fields are in character format.

Expansion level 1 is 254 bytes long and contains fields 1 through 30. If you specify expansion level 2 in the Audit Browse command or the Audit Retrieve command, the message audit record is 326 bytes long and contains fields 1 through 42. If you specify expansion level 3 in the Audit Browse command or the Audit Retrieve command, the message audit record is 340 bytes long and contains fields 1 through 43. If you specify expansion level 4 in the Audit Browse command or the Audit Retrieve command, the message audit record is 403 bytes long and contains fields 1 through 53.

#	COL	SIZE	NAME	DESCRIPTION
1	1	16	AUSERID	Account ID and user ID of record owner
2	17	1	ARECTYP	Type of audit record
3	18	20	APARTNER	Trading partner ID
4	38	18	AGROUPID	Internal Information Exchange message group identifier
5	56	1	ASTATFLG	Status flag
6	57	15	ASTATUS	Status phrase
7	72	3	APURGFLG	Purge reason code
8	75	15	APURGED	Purge phrase
9	90	20	AALIAS	Alias reference to trading partner
10	110	1	APRIORITY	Message priority
11	111	1	ANMSGCLS	Network message class
12	112	8	AUMSGCLS	User message class
13	120	8	AMSGNAME	User message name
14	128	5	AMSGSEQI	Message sequence
15	133	8	ASYSTYP	Sender's system type
16	141	4	ASYSLVL	Sender's system level
17	145	8	AMSGCNTR	Number of messages in message group
18	153	8	AMSGSIZE	Total text size
19	161	8	ARCVARCH	Receiver's archive ID for message
20	169	6	ASNDDATE	Date sent (queued/forwarded/mailboxed) (YYMMDD)
21	175	6	ASNDDTIME	Time sent (queued/forwarded/mailboxed) (HHMMSS)
22	181	6	ARCVDATE	Date purged/delivered (YYMMDD)
23	187	6	ARCVDTIME	Time purged/delivered (HHMMSS)
24	193	6	AMSEQOUT	Message output sequence number
25	199	1	AOETYPE	EDI data type
26	200	17	AOESEND	EDI sender ID
27	217	4	AOESQUAL	EDI sender ID qualifier

#	COL	SIZE	NAME	DESCRIPTION
28	221	17	AOERCVR	EDI receiver ID
29	238	4	AOERQUAL	EDI receiver ID qualifier
30	242	13	AOEMSGID	EDI message ID

Expansion Level 2—Extended timing fields 31 through 42:

#	COL	SIZE	NAME	DESCRIPTION
31	255	6	AOSNDEDD	Date send ended (YYMMDD)
32	261	6	AOSNDEDT	Time send ended (HHMMSS)
33	267	6	AOSNDCMD	Date send committed (YYMMDD)
34	273	6	AOSNDCMT	Time send committed (HHMMSS)
35	279	6	AOSNDROD	Date routed (YYMMDD)
36	285	6	AOSNDROT	Time routed (HHMMSS)
37	291	6	AORCVSTD	Date receive started (YYMMDD)
38	297	6	AORCVSTT	Time receive started (HHMMSS)
39	303	6	AORCVEDD	Date receive ended (YYMMDD)
40	309	6	AORCVEDT	Time receive ended (HHMMSS)
41	315	6	AORCVCMD	Date receive committed (YYMMDD)
42	321	6	AORVCVCT	Time receive committed (HHMMSS)

Expansion Level 3—EDI exchange control number field 43:

#	COL	SIZE	NAME	DESCRIPTION
43	327	14	AOEDICTL	EDI exchange control number

Expansion Level 4—Carbon copy and other miscellaneous fields; fields 44 through 53:

#	COL	SIZE	NAME	DESCRIPTION
44	341	8	AOCHILDID	Cluster child ID
45	349	1	AOCHARGE	Resolved message charge
46	350	9	AOMAILBAG	Mailbag number for VANS interconnect
47	359	10	AORMSIZE	Size of message received
48	369	10	AOSMSIZE	Size of message sent
49	379	1	AOPAYPRC	Who pays carbon copy processing charge
50	380	1	AOPAYRCV	Who pays carbon copy receive charge
51	381	1	AOMSGTYP	Message type
52	382	1	AO3RDPTY	Third party ID type
53	383	21	AO3RDPRT	Third party ID

AUSERID (field 1)

This field contains the Information Exchange address of the following user:

- For send audits, this field represents the sender of the message.
- For receive audits, this field represents the receiver of the message.
- For carbon copy as partner audits, this field represents the partner in the carbon copy relationship.

The first 8 characters represent the user's Information Exchange account ID; the last 8 characters are the user ID.

ARECTYP (field 2)

This field contains one of the following numeric values:

This value: Indicates:

- | | |
|---|--|
| 0 | An audit record for a sent message |
| 1 | An audit record for a message sent to yourself |
| 2 | An audit record for a received or receivable message |
| 3 | An audit record for Carbon Copy as Partner |

APARTNER (field 3)

This field contains the Information Exchange address of the following user:

- For send audits, this field represents the receiver of the message.
- For receive audits, this field represents the sender of the message.
- For carbon copy as partner audits, this field represents the requester of the carbon copy message.

If the trading partner and the user specified in the `AUSERID` field are on the same system, the first 4 characters are blanks. If not, the first character is *I*, followed by the 3-character system ID of the trading partner's system. The next 8 characters represent the user's Information Exchange account ID; the last 8 characters are the user ID.

AGROUPID (field 4)

This field contains the ID identifier for the message group. Each message group is assigned an internal control number. This field (together with the `APARTNER` and `AUSERID` fields) uniquely identifies the message and message recipient. If used alone, `AGROUPID` does not ensure a unique database key (if, for example, a message is sent to a list of users).

ASTATFLG (field 5)

This field contains a value that determines the status of the message group.

This value:	Indicates:
0	The message is being transferred to another system. The transfer is not complete.
1	The message is received by the trading partner, or the message is transferred to a non-Information Exchange system.
2	The message is purged and the reason is displayed in APURGFLG (field 7) of the audit record.
3	The message is in the trading partner's mailbox.
4	The message is queued to another value added network (VAN). The transfer process is not complete.
5	The message is transferred to another system. The other system has accepted the message but does not support message disposition updates. The final disposition (Received/Purged) of the message is unknown. An example is to another VAN that does not support Interchange Delivery Notice (TA3).
6	The message is forwarded to another system. The other system has accepted the message and will inform Information Exchange when: <ul style="list-style-type: none"> • The message is placed in the trading partner's mailbox • The message is received • The message is purged
7	The message has been redirected to another mailbox as a result of carbon copy processing. This is the final status for this message.

The following table shows the status changes for messages sent to a trading partner on a non-X.400 system.

Messages sent to a trading partner on the following:	Can have these status changes:
The local Information Exchange system	Mailboxed (3) ®Received (1) or Purged (2) or Redirected (7)
Another Information Exchange system, the Information Exchange Peer System, and the EDI Server	In Transit (0) ®Mailboxed (3) ®Received (1) or Purged (2)
A non-Information Exchange system (such as IBM Mail Exchange)	In Transit (0) ®Sent to non-Information Exchange (1)
A VAN that supports TA3 and messages to IBM MQSeries Services users	In Transit (0) ®Queued (4) ®Forwarded (6) ®Mailboxed (3) ®Received (1) or Purged (2)
A VAN that does not support TA3	In Transit (0) ®Queued (4) ®Forwarded (6)

The following table shows the status changes for P2 messages sent to a trading partner on an X.400 system.

P2 messages sent to a trading partner on the following:	Can have these status changes:
An X.400 system that supports Delivery Reports (DR) and Non-Delivery Reports (NDR)	In Transit (0) ®Forwarded (6) ®Received (1) or Purged (2)
An X.400 system that does not support DR or NDR	In Transit (0) ®Forwarded (6) ®Transferred (5)

The following table shows the status changes for P35 messages sent to a trading partner on an X.400 system.

P35 messages sent to a trading partner on the following:	Can have these status changes:
An X.400 system that supports DR, NDR, and EDI Notification (EDIN)	In transit (0) ®Forwarded (6) ®Mailboxed (3) ®Received (1) or Purged (2)
An X.400 system that does not support DR or NDR, but does support EDIN	In transit (0) ®Forwarded (6) ®Received (1) or Purged (2)
An X.400 system that supports DR and NDR, but does not support EDIN	In transit (0) ®Forwarded (6) ®Mailboxed (3) ®Transferred (5)
An X.400 system that does not support DR, NDR, or EDIN	In transit (0) ®Forwarded (6) ®Transferred (5)

ASTATUS (field 6)

This field contains a short-phrase indication of one of the above status values.

APURGFLG (field 7)

This field contains a value that gives more information on the reason a message was purged. If a message has been purged, you can determine the reason by looking at the value in APURGFLG (field 7) of the audit record. Information Exchange purges a message as a result of a Cancel command by the sender or intended recipient, or when certain system-related events occur. The following lists the values you may find in APURGFLG:

This value:	Indicates:
001	The message was purged from the user's mailbox because the retention period for the message expired.
002	The message was canceled by the sender.
003	The message was canceled by the intended receiver.
004	An administrator canceled the message.
005	The message was canceled by network personnel.
008	The message was canceled by X, where X is assigned by the user's local system.
010	The message is not deliverable.
011	The message cannot be delivered immediately.
016	The message is processed by the Information Exchange Administration mailbox.
017	The receiver's user ID was deleted from Information Exchange.

APURGED (field 8)

This field contains a short-phrase indication of one of the status values shown above, or if the system did not purge the message group, this field contains blanks.

AALIAS (field 9)

This field contains an alias if the sender of the message group is AUSERID and the message was sent to an alias.

APRIORITY (field 10)

This field contains the message priority as sent on the Send Message command.

ANMSGCLS (field 11)

This field contains the network message class of the message group.

AUMSGCLS (field 12)

This field contains the user message class of the message group.

AMSGNAME (field 13)

This field contains the name of the message group.

AMSGSEQI (field 14)

This field contains the sender's assigned message sequence.

ASYSTYP (field 15)

This field contains the message sender's system type.

ASYSLVL (field 16)

This field contains the message sender's system level.

AMSGCNTR (field 17)

This field contains the number of messages in the message group, excluding any common data header message.

AMSGSIZE (field 18)

This field contains the total text size of the message group, excluding any Information Exchange headers or common data headers. This field and AOSMSIZE (field 48) may contain the same information; however, this field allows 8 bytes to be returned and AOSMSIZE allows 10 bytes.

ARCVARCH (field 19)

This field contains the receiver's archive ID or the session key of the session where it was received. If the status is not received, this field contains blanks.

ASNDDATE (field 20)

This field contains the date when the message was sent, offset to the time zone specified in either the Audit Browse or Audit Retrieve command. The format is *YYMMDD*.

ASNDDATE (field 21)

This field contains the time when the message was sent, offset to the time zone specified in either the Audit Browse or Audit Retrieve command. The format is *HHMMSS*.

ARCVDATE (field 22)

This field contains the date when the message group was purged or delivered, offset to the time zone specified in either the Audit Browse or Audit Retrieve command. The format is *YYMMDD*. If the status is not purged or delivered, this field contains zeros.

ARCVTIME (field 23)

This field contains the time when the message group was purged or delivered, offset to the time zone specified in either the Audit Browse or Audit Retrieve command. The format is *HHMMSS*. If the status is not purged or delivered, this field contains zeros.

AMSEQOUT (field 24)

This field contains the sequence-out number that Information Exchange assigns to the first message of the message group when it is received.

AOETYPE (field 25)

This field contains a value that indicates the EDI data type.

This value:	Indicates:
0	Information Exchange does not recognize this as one of the supported EDI data types.
1	<i>American National Standards Institute (ANSI) X12 data.</i>
2	<i>Uniform Communications Standard (UCS) data.</i>
3	<i>Electronic Data Interchange for Administration, Commerce, and Transportation (EDIFACT) data.</i>
4	<i>United Nations/Trade Data Interchange (UN/TDI) data.</i>

AOESEND (field 26)

This field contains the EDI sender ID and is only for EDI use. The format is:

This EDI data type:	Requires this format:
X12	15 alphanumeric characters, left-justified, and padded on the right with blanks.
UCS	2 to 12 alphanumeric characters, left-justified, and padded on the right with blanks.
EDIFACT	Up to 17 characters of the data element 0004 in composite data element S002 (interchange Sender).
UN/TDI	If the first subelement of the sender ID element (FROM:1) is not blank, then it is used as the EDI send ID. Otherwise, up to 17 characters of the second subelement of the sender ID (FROM:2) are used.

AOESQUAL (field 27)

This field contains a qualifier that gives more information about the EDI sender. The format is:

This EDI data type:	Requires this format:
X12	2 alphanumeric characters, left-justified, and padded on the right with blanks.
UCS	No qualifier used; the field is blank.
EDIFACT	4 alphanumeric characters.
UN/TDI	No qualifier user ID; the field is blank.

AOERCVR (field 28)

This field is only for EDI use and contains the EDI receiver ID. The format is shown below:

This EDI data type:	Requires this format:
X12	15 alphanumeric characters, left-justified, and padded on the right with blanks.
UCS	2 to 12 alphanumeric characters, left-justified, and padded on the right with blanks.
EDIFACT	Up to 17 characters of the data element 0010 in composite data element S003 (Interchange Recipient).
UN/TDI	If the first subelement of the receiver ID element (UNTO:1) is not blank, then it is used as the EDI receiver ID. Otherwise, up to 17 characters of the second subelement of the receiver ID (UNTO:2) are used.

AOERQUAL (field 29)

This field contains a qualifier that gives more information about the EDI receiver. The format is show below:

This EDI data type:	Requires this format:
X12	2 alphanumeric characters, left-justified, and padded on the right with blanks.
UCS	No qualifier used; the field is blank.
EDIFACT	4 alphanumeric characters.
UN/TDI	No qualifier userid; the field is blank.

AOEMSGID (field 30)

For non-EDI data, this field contains the message ID, which has two elements: the message name (8 bytes) and the sequence number (5 bytes). For EDI data, the format is shown below:

This EDI data type:	Requires this format:
X12	The 9-digit interchange control number, right-justified, and padded on the left with zeros.
UCS	The 1- to 5-digit interchange control number, right-justified, and padded on the left with zeros.
EDIFACT	Up to 13 characters of the sender's data element 0020 (Interchange Control Reference). This field is left-justified and padded on the right with blanks.
UN/TDI	Up to 13 characters of the sender's reference element (SNRF). This field is left-justified and padded on the right with blanks.

For expansion levels greater than 1, the following fields are returned (all these data time fields are offset to the time zone specified in either the Audit Browse or Audit Retrieve command):

AOSNDEDD (field 31)

This field contains the date when the send process ended. The format is *YYMMDD*.

AOSNDEDT (field 32)

This field contains the time when the send process ended. The format is *HHMMSS*.

AOSNDCMD (field 33)

This field contains the date when Information Exchange committed the message during the send process. The format is *YYMMDD*.

AOSNDCMT (field 34)

This field contains the time when Information Exchange committed the message during the send process. The format is *HHMMSS*.

AOSNDROD (field 35)

This field contains the date when the message was queued, forwarded, or mailboxed. The format is *YYMMDD*.

AOSNDROT (field 36)

This field contains the time when the message was queued, forwarded, or mailboxed. The format is *YYMMDD*.

AORCVSTD (field 37)

This field contains the date when the purge or delivery process started. The format is *YYMMDD*.

AORCVSTT (field 38)

This field contains the time when the purge or delivery process started. The format is *HHMMSS*.

AORCVEDD (field 39)

This field contains the date when the purge or delivery process ended. The format is *YYMMDD*.

AORCVEDT (field 40)

This field contains the time when the purge or delivery process ended. The format is *HHMMSS*.

AORCVCMD (field 41)

This field contains the date when the purge or delivery process was committed. The format is *YYMMDD*.

AORCVCMT (field 42)

This field contains the time when the purge or delivery process was committed. The format is *HHMMSS*.

For expansion levels greater than 2, the following fields are returned:

AOEDICTL (field 43)

This field contains the EDI exchange control number used to identify the EDI exchange. For non-EDI data, this field is blank. The EDI exchange control number is assigned as follows per each standard listed:

This EDI data type: Requires this format:

X12	The 9-digit interchange control number, right-justified, and padded on the left with zeros.
UCS	The 1- to 5-digit interchange control number, right-justified, and padded on the left with zeros.
EDIFACT	Up to 14 characters of the sender's data element <i>0020</i> (Interchange Control Reference). This field is left-justified and padded on the right with blanks.
UN/TDI	Up to 14 characters of the sender's reference element (SNRF). This field is left-justified and padded on the right with blanks.

For expansion levels greater than 3, the following fields are returned:

AOCHILDID (field 44)

This field contains the cluster child ID if the message was sent or received by a cluster child. Otherwise, blanks are returned.

AOCHARGE (field 45)

This field contains the charge mode that was actually used by Information Exchange. One of the following values is returned: **1** for receiver pays all charges, **6** for sender pays all charges, or **5** for split charges. Intransit messages and carbon copy/redirect messages return blanks in this field.

AOMAILBAG (field 46)

This field contains the mailbag number assigned to the message when sent to a VAN interconnected user. If the message is not sent through the VAN interconnect service, blanks are returned.

AORMSIZE (field 47)

This field contains the total size of the message group received, excluding any Information Exchange headers or common data header.

AOSMSIZE (field 48)

This field contains the total size of the message group sent, excluding any Information Exchange headers or common data headers. This field and AMMSGSIZE (field 18) may contain the same information; however, this field allows 10 bytes to be returned and AMMSGSIZE allows 8 bytes.

AOPAYPRC (field 49)

This field contains information as to who pays for the carbon copy *processing* charge. A value of 0, 1, 2, or 3 may be returned in this field.

This value: Indicates who pays carbon copy *processing* charge:

0 This is not a carbon copy/redirect message.

1 Sender pays charges.

2 Receiver pays charges.

3 Partner pays charges.

AOPAYRCV (field 50)

This field contains information as to who pays for the carbon copy *receive* charge. A value of 0, 1, 2, or 3 may be returned in this field.

This value: Indicates who pays carbon copy *receive* charge:

0 This is not a carbon copy/redirect message.

1 Sender pays charges.

2 Receiver pays charges.

3 Partner pays charges.

AOMSGTYP (field 51)

This field contains the message type. A value of 0, 1, 2, 3, or 4 may be returned in this field.

This value:	Indicates the message type:
0	This is not a carbon copy/redirect message.
1	This is a send requested carbon copy.
2	This is a receive requested carbon copy.
3	This is a receive requested redirect.
4	This is an original message that was redirected.

AO3RDPTY (field 52)

This field contains the third party ID type. A value of 0, 1, 2, or 3 may display in this field. The value in this field is used to determine the format of the AO3RDPRT (field 53).

This value:	Indicates:
0	Blanks are returned.
1	An EDI reference ID consisting of: <ul style="list-style-type: none">• A 4-byte EDI qual ID, left-justified and padded on the right with blanks. If not applicable, this field is blank.• A 17-byte EDI ID, left-justified and padded on the right with blanks.
2	An alias reference ID consisting of a 1-byte table type, a 3-byte table name, and a 16-byte alias name, padded on the right with one blank.
3	A native Information Exchange ID consisting of a 3-byte system ID (or blanks for local users), an 8-byte account ID, and an 8-byte user ID, padded on the right with 2 blanks.

AO3RDPRT (field 53)

This field contains the third party ID in the form indicated by the A03RDPTY (field 52) associated with carbon copy/redirect messages. The content depends on the values in ARECTYP (field 2) and AOMSGTYP (field 51).

When ARECTYPE is:	When AOMSGTYP is:	A03RDPRT contains:
0 or 1 (Send audit)	0 (not carbon copy) 1 (send carbon copy) 2 (receive carbon copy) 3 (redirect) 4 (original redirected)	Blanks Original receiver Original sender Original sender Blanks
1 or 2 (Receive audit)	0 (not carbon copy) 1 (send carbon copy) 2 (receive carbon copy) 3 (redirect) 4 (original redirected)	Blanks Original receiver Original sender Original sender Carbon copy recipient
3 (Carbon copy as partner audit)	1 (send carbon copy) 2 (receive carbon copy) 3 (redirect)	Carbon copy recipient Carbon copy recipient Carbon copy recipient

Archive commands

Archive Retrieve command—SDIARTV

This command causes Information Exchange to move messages identified by the archive reference ID from archive storage to your normal-priority message queue. You must receive these messages using the Receive Message command. The length of the Archive Retrieve command is 49 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	SDIARTV
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	8	RETRSPC	Response command name
7	42	8	ARREFID	Archive reference ID

COMMAND (field 1)**ACCNTNO (field 2)****USERID (field 3)****SESSKEY (field 4)****EXPAND (field 5)**

This field contains a command expansion-level indicator. It must contain the value 1. Invalid values default to 1.

RETRSPC (field 6)

This field contains the value that is returned in the first 8 characters of the Archive Retrieve response. The format is alphanumeric, left-justified, and padded on the right with blanks.

ARREFID (field 7)

This field specifies the archive reference identifier of the messages you want to retrieve from short-term archive storage. Use the SDIRCVM command to receive them. The format is alphanumeric, left-justified, and padded on the right with blanks.

Archive Retrieve response

This is Information Exchange's response to the Archive Retrieve command; its length is 38 bytes.

#	COL	SIZE	NAME	DESCRIPTION
1	1	8	COMMAND	Value from RETRSPC of SDIARTV
2	9	8	ACCNTNO	Account ID
3	17	8	USERID	User ID
4	25	8	SESSKEY	Session access key
5	33	1	EXPAND	Expansion indicator
6	34	5	RETVCNT	Number of archive messages retrieved

COMMAND (field 1)

ACCNTNO (field 2)

USERID (field 3)

SESSKEY (field 4)

EXPAND (field 5)

This field contains the value from EXPAND (field 5) of the Archive Retrieve command.

RETVCNT (field 6)

This field contains the number of archived message groups the Archive Retrieve command retrieves and sends to your normal-priority message queue. If this field contains zero, either an unknown archive reference was used in the Archive Retrieve command, or all the messages from that archive group have already been retrieved and are still in the mailbox. The format is numeric, right-justified, and padded on the left with zeros.



Glossary

This glossary defines words as they are used in this book. It includes terms and definitions from the Dictionary of Computing (SC20-1699). If you are looking for a term and cannot find it here, reference this book for additional definitions.

A

account. (1) A set of users who work in a company. (2) A set of users who work in a related area of a company.

account ID. (1) The network's designation for the group of users with whom you are associated. (2) A name that identifies an account to Information Exchange.

acknowledgment. Response from Information Exchange that lets you know whether your messages were delivered, received, purged, or various combinations.

address. A user's account name and user identification (ID); used by Information Exchange to route messages.

alias. A nickname used in place of an account ID and user ID (Information Exchange address).

alias table. A nickname file kept within Information Exchange.

alphanumeric. Generally, any keyboard character, but for practical purposes should be restricted to alphabetic, numeric, space, and common punctuation characters.

American National Standards Code for Information Interchange (ASCII). The standard code using a coded character set consisting of 7-bit coded characters (8 bits including parity check) used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters.

American National Standards Institute (ANSI). An organization consisting of producers, consumers, and general interest groups that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States.

ANSI. American National Standards Institute.

ANSI X12. EDI standard for transmission of business data.

archive. A place in Information Exchange to store messages you receive from Information Exchange.

ASCII. American National Standard Code for Information Interchange.

asynchronous. A protocol that enables a communication device to operate in an unsynchronized and unpredictable manner with respect to the start of program instructions. This type of communication is used for modems and low-speed ASCII terminals (PCs); see “protocol.”

attribute. A property or characteristic of one or more entities (for example, length, value, color, or intensity).

audit. A way of tracking and verifying basic information about the status of messages.

C

carriage-return and line-feed (CRLF) characters. A word processing formatting control that moves the printing or display point to the first position of the next line.

CDH. Common data header.

character. Generally, any letter, number, punctuation mark, or other symbol that can be entered from a computer keyboard.

command. A character or group of characters you enter that tells the system to perform some task.

commit. The point at which a message is either delivered, canceled, or purged. When a session fails, all uncommitted messages are lost.

common data header. Used by Information Exchange interfaces to provide an additional description of a message group.

CRLF. Carriage-return and line-feed characters.

D

default. A value an application program automatically uses until you specify a different value.

delivery class. Specifies how messages and files are delivered; senders can choose from high-priority, normal-priority, and express-priority delivery.

destination. The final Information Exchange address to which a message is to be delivered.

distribution list. A list of the addresses of users with whom a certain user communicates; used to send messages to several people without having to type their addresses.

E

EBCDIC. Extended binary-coded decimal interchange code.

EDI. Electronic data interchange.

EDIFACT. Electronic data interchange for administration, commerce, and transportation.

Electronic Data Interchange (EDI). The process of sending specially formatted business documents directly from one computer to another electronically.

Electronic Data Interchange For Administration, Commerce, And Transportation (EDIFACT). An EDI standard for the fields of administration, commerce, and transportation.

electronic mail. The process of sending free-formatted messages from one computer to another.

envelope. Information Exchange's electronic envelope contains the text of a message, together with a message header. See also “message header.”

ESO. Extended Security Option.

Extended Binary-coded Decimal Interchange Code (EBCDIC). A coded character set consisting of 256 8-bit coded characters.

Extended Security Option (ESO). An option that can be specified on the user profile, for users who require stricter security of data (extended security users).

F

field. An area of a command reserved for data of a certain type or length.

file. A document that you can send or receive using your interface.

G

global alias. An alias that can be used by any Information Exchange user on a particular system.

global alias table. (1) A system-wide alias table. (2) An alias table set up within a system.

H

header. See “message header.”

high-priority message. Messages that move to the front of the queue when they are received into Information Exchange; normal-priority messages enter the queue in the order that they are received into Information Exchange.

I

Information Exchange. A communication service that enables users to send and receive information electronically to and from trading partners.

Information Exchange address. A combination of a user's account ID and user ID; used to route messages.

Information Exchange mailbox. Information Exchange's electronic mailbox. See also “electronic mailbox.”

L

library. A place to store data for an extended period of time.

M

mailbox. See “electronic mailbox.”

member. The text of a message or document that is stored in a library.

message. (1) Any piece of data that users send or receive. (2) The smallest subdivision of information that can be sent from one user to another. (3) An instruction or explanation on the screen that tells you what the system is doing or warns you that the system has detected an error.

message acknowledgment. A response from Information Exchange that lets users know whether their messages were delivered, received, purged, or various combinations of the three.

message class. A category used to group mail, agreed on among trading partners.

message header. The beginning of a message, prior to the text; it serves as both the command and the address. When a message is sent from Information Exchange, the header indicates the addresses of both the sender and the intended recipient.

message sequence. The order of messages within a group of messages.

O

organizational alias. (1) An alias that can be used by any user in an account. (2) A company-wide alias table.

organizational alias table. An alias table set up within an account.

owning account. The account that owns a library.

P

parameter. (1) A variable that is given a constant value for a specified application and that may denote the application. (2) An item in a menu for which you specify a value or for which the system provides a value when the menu is interpreted. (3) The data passed between programs or procedures.

password. A combination of characters you enter when you log on, designed to be secret to prevent unauthorized access of your data.

peer system. An Information Exchange peer system is another EDI system for customers requiring special processing options

private alias. (1) An alias that can be used only by the user who created it. (2) An individual user's alias table.

private alias table. An alias table set up for an individual user.

protocol. A prearranged procedure for communication.

R

receiver. The user or users to whose mailboxes you are sending, or retrieving information.

response. Information Exchange's reply to certain commands (known as response-mode commands).

S

session. The period of time during which you can communicate with a computer system or one of its programs; usually, the elapsed time between logon and logoff.

SNA. Systems Network Architecture.

system ID. A unique identifier assigned to a unique Information Exchange system.

Systems Network Architecture (SNA). The description of the logical structure, formats, protocols, and operating sequences for transmitting information units through, and controlling the configuration and operation of, networks.

T

trading partner. A business associate with whom you exchange information electronically.

U

UCS. Uniform Communication Standard.

Uniform Communication Standard (UCS). A standard EDI format used in the grocery industry.

United Nations/trade Data Interchange (UN/TDI). An EDI standard for administration, commerce, and transportation fields developed by the United Nations Economic Commission for Europe.

UN/TDI. United Nations/Trade Data Interchange.

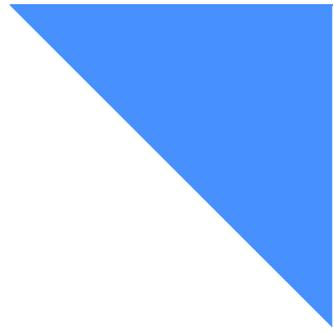
user classification. A short description that users can assign to their documents to identify their documents to their trading partners.

user identification (user ID). A name that identifies a user to Information Exchange within an account.

user profile. A list of the characteristics of how a user works with Information Exchange.

X

X12. A specially formatted data stream.



Index

A

acknowledgments
 DELIVERY 39
 PURGE 39
 RECEIPT 39
 requesting 39
 types 39
adding members to the library 62
address
 Information Exchange 12
addresses
 Information Exchange 45
 intersystem Information Exchange 45
 pseudonyms 47
addressing messages 45
alias addressing 47
Alias Inquiry
 command format 176
 response format 178
Archive Retrieve
 command format 258
 response format 259
Archive Retrieve response 259
archive storage days 8
Archiving carbon copy messages 60
archiving messages 59
Asynchronous background processing 42
asynchronous background processing 42
Asynchronous Probe response message text 129
audit
 data retrieval 41
 storage 8

 trail 40
Audit Browse
 command format 231
 response format 236
Audit Browse command 41, 231
Audit Record/Display format 243
Audit Retrieve
 command format 237
 response format 242
Audit Retrieve command 41, 237
Audit Retrieve response 242

B

Browse Library Member
 command format 190
 response format 192
Browse Library Member command 190
Browse Library Member response 192
browsing library members 63

C

Cancel
 command format 118
Cancel command 38
 format 118
canceling messages 38
carbon copy
 cluster userids 79
 non-candidates for 31
 processing charge 5
 receive charge 5

- requesting messages 29
- Carbon copy function 30
- CDH
 - creating 24
 - entry fields 69
 - errors 25
 - format 24
 - receiving 24
 - Retrieve Library Member 219
 - Retrieve Library Member response 220
- CDH (common data header) 219
- CDH field
 - FILETIME 74
 - LIBMEMB 75
 - UNIQUEID 73
- CDH fields
 - CODEPAGE 73
 - COMFNAME 76
 - COMPDLN 76
 - COMPSOFT 76
 - COMPSVER 76
 - CRLFEOF 74
 - CTYPE 71
 - DELIMIT 72
 - DESCRIBE 72
 - DIFORMAT 72
 - EDICNTLN 75
 - EDIRECVR 75
 - EDIRQUAL 74
 - EDISENDER 74
 - EDISQUAL 74
 - FILEDATE 74
 - FILENAME 71
 - LIBACCT 75
 - LIBNAME 75
 - LIBREPL 76
 - LOCATION 71
 - ORGLN 75
 - RECORDF 71
 - RESERVED 74
 - RLENGTH 71
 - SYSTYPE 73
 - TTYPE 72
 - VERSION 74
- character format
 - Session Trace Receive Record 111
 - Session Trace Record 115
- character-code sets 50
- checkpoint-level recovery 18
- cluster mailbox 6
- cluster userids
 - carbon copy 79
- commands
 - causing responses 81
 - expansion of 85
 - format of 85
 - Alias Inquiry 176
 - Archive Retrieve 258
 - Audit Browse 231
 - Audit Retrieve 237
 - Browse Library Member 190
 - Cancel 118
 - Commit 123
 - Define Alias 179
 - Define Library 194
 - Delete Library 199
 - Delete Library Member 201
 - Library Title 227
 - List Define 182
 - List Libraries 203
 - List Library Members 210
 - List Verify 186
 - Load Test Messages 144
 - Message Inquiry 146
 - Message Queue Query 149
 - Pace 125
 - Probe 127
 - Purge Input 133
 - Purge Message 159
 - Receive Message 161
 - Retrieve Library Member 214
 - Retrieve Library Member CDH 217, 219
 - Search Library 222
 - Send Message 134
 - Session End 96
 - Session Inquiry 99
 - Session Start 87
 - Session Trace Browse 104
 - Session Trace Retrieve 107
- names
 - SDIARTV 258
 - SDIAUDR 237
 - SDICMIT 123
 - SDICNCL 118
 - SDIDALS 179
 - SDIINQA 176
 - SDIINQM 146
 - SDIINQS 99
 - SDILBRW 190
 - SDILDEF 194
 - SDILDEL 199
 - SDILDLM 201
 - SDILLST 203
 - SDILMBR 210

- SDILRTV 214
 - SDILSRCH 222
 - SDILSTD 182
 - SDILSTV 186
 - SDILTST 144
 - SDIPRGM 159
 - SDIPROB 127
 - SDIQUMS 149
 - SDIRCDH 219
 - SDIRCVM 161
 - SDISEND 96
 - SDISNDM 134
 - SDISSTA 87
 - SDITITL 227
 - SDITRLB 104
 - SDITRLR 107
 - overview 4
 - Receive Message 55
 - Retrieve Library Member CDH 219
 - Commit command 51, 123
 - Commit request 142
 - format of 142
 - frequency 56
 - using 56
 - Commit response 52, 57
 - committing messages 51, 52
 - common data header 69
 - creating 24
 - errors 25
 - format 24
 - receiving 24
 - communication protocols 2
- D**
- data
 - audit retrieval 41
 - meaningful exchange 50
 - organizing for transmission efficiency 49
 - self-defining 50
 - separating
 - by destination 49
 - by processing requirements 49
 - by type 49
 - data-element splitting 50
 - Define Alias command 47, 179
 - Define Library command 194
 - defining a library 62
 - Delete Library command 199
 - Delete Library Member command 201
 - Delete Library Member response 202
 - Delete Library response 200
 - deleting a library 62
 - deleting members from the library 62
 - distribution lists
 - permanent 47
 - temporary 47
 - verifying 48
- E**
- ending a session 14
 - extended security option 11
 - Extended Security Option (ESO) 11
 - password 11
 - user ID 12
- G**
- group-level recovery 20
- H**
- headers
 - long format 26
 - short format 26
- I**
- Information Exchange
 - addresses 12, 45
 - features 6
 - service parameters 8
 - session 11
 - Intersystem Information Exchange addresses 45
 - item
 - format
 - Library Member 230
 - structure
 - Library Entry 207
- L**
- library
 - adding members 62
 - browsing library members 63
 - commands
 - Browse Library Member 63
 - Define Library 62
 - Delete Library 62
 - List Libraries 64
 - List Library Members 64
 - Retrieve Library Member 63
 - description 61

- listing libraries 64
- listing members 64
- members 61
- owner 61
- retrieving a library member 63
- searching the library 64
- updating library members 63
- Library Entry item structure 207
- Library Member
 - item format 230
 - list entry format 213
- Library Member List entry 213
- Library Title command 227
- Library Title response 228
- List Define command 47, 182
- List Libraries command 203
- List Libraries response 205
- List Library Members command 210
- List Library Memembers response 211
- list searches 47
- List Verify command 48, 186
- Load Test Messages command 39, 144
- long format, Received Message header 169, 170
- long-format headers 26

M

- message
 - addressing 45
 - affects on receiving 13
 - archiving 59
 - audit feature 41
 - canceling 38
 - carbon copy 7
 - charges 5
 - committal 15
 - creating archive groups 59
 - delivery classes 6
 - envelope 23
 - group 24, 27
 - defining 28
 - headers 23
 - inquiry 40
 - large 25
 - limits 9
 - long transactions 27
 - multiple-segment 26
 - physical block size 25
 - protecting 51
 - receiving 19
 - record delimiters 23
 - restarting transmission 19

- retention days 9
- retrieving from archive 59
- selecting and receiving test 39
- sending and delivering correctly 39
- transmitting 18
- type 23
- using segments 50
- message and session tracking 40
- message charges
 - carbon copy 5
- Message delivery classes 44
- Message Inquiry command 146
- Message Inquiry response 147
- Message Queue entry 153
- Message Queue Query command 38, 149
- Message Queue Query response 151
- message segments 25
- message-grouping facility 27
- message-level recovery 20
- messages
 - receiving 6
 - sending 6

N

- network security 3

P

- Pace command 52, 125
- Pace request 57, 157
 - format of 157
 - responding to 57
- Pace response 53
- Pace response (received) 126
- Pace response (sent) 158
- Probe command 49, 127
- Purge Input command 133
- Purge Message command 159
- Purge Message response 160

R

- Receive Message command 55, 161
- Received Message header
 - long format 169, 170
 - short format 171
- receiving messages 13, 19
- recovery
 - checkpoint-level 18
 - group-level 20
 - levels 16

- message-level 20
 - session-level 16
 - stream-level 17
 - recovery and restart 14
 - recovery point 51
 - Redirect function 29
 - requesting acknowledgments 39
 - requesting test messages 39
 - requests 4
 - response
 - format of
 - Alias Inquiry 178
 - Archive Retrieve 259
 - Audit Browse 236
 - Audit Retrieve 242
 - Browse Library Member 192
 - Commit (sent) 143
 - Delete Library 200
 - Delete Library Member response 202
 - Library Title 228
 - List Libraries 205
 - List Library Members 211
 - Message Inquiry 147
 - Message Queue Query 151
 - Pace (received) 126
 - Pace (sent) 158
 - Purge Message 160
 - Retrieve Library Member 217
 - Retrieve Library Member CDH 220
 - Search Library 224
 - Session End 97
 - Session Inquiry Response 101
 - Session Inquiry response 100
 - Session Start 91
 - Session Trace Browse response 106
 - Session Trace Retrieve 110
 - Synchronous Probe 131
 - overview 4
 - response codes
 - Session End 98
 - Session Start 93
 - restart and recovery 14
 - restarting a receiving operation 20
 - restarting transmission 19
 - Retrieve Library Member command 214
 - Retrieve Library Member response 217
 - retrieving a library member 63
- S**
- SDIARTV 258
 - SDIAUDR 237
 - SDICMIT 123
 - SDICNCL 118
 - SDICRSP 143
 - SDIDALS 179
 - SDIINQA 176
 - SDIINQM 146
 - SDIINQS 99
 - SDILBRW 190
 - SDILDEF 194
 - SDILDEL 199
 - SDILDLM 201
 - SDILLST 203
 - SDILMBR 210
 - SDILRTV 214
 - SDILSRCH 222
 - SDILSTD 182
 - SDILSTV 186
 - SDILTST 144
 - SDIPRGM 159
 - SDIPROB 127
 - SDIPRSP 158
 - SDIQUMS 149
 - SDIRCDH 219
 - SDIRCVM 161
 - SDISEND 96
 - SDISNDM 134
 - SDISSTA 87
 - SDITITL 227
 - SDITRLB 104
 - SDITRLR 107
 - Search Library command 222
 - Search Library response 224
 - searching the library 64
 - security, network 3
 - Self-defining data 50
 - Send Message command 43, 49, 134
 - long format 134
 - short format 135
 - session
 - abnormal end 14
 - defined 3
 - Session End command 14, 96
 - Session End response 14, 97
 - session inquiry 40
 - Session Inquiry command 99
 - Session Inquiry response
 - long format 100
 - short format 101
 - Session Start command 12, 87, 91
 - Session Start response 14
 - session trace 40
 - Session Trace Browse command 42, 104

- Session Trace Browse response 106
- session trace feature 42
- Session Trace Receive Record--character format 111
- Session Trace Record--character format 111
- Session Trace Retrieve command 42, 107
- Session Trace Retrieve response 110
- session-level recovery 16
- short format, Received Message header 171
- short-format headers 26
- SNA protocols 20
- space compression
 - preventing 51
 - using asterisks 51
- starting a session 12
- status information 40
- storage time limit 40
- stream-level recovery 17
- supported communication protocols 2
- Synchronous Probe response 131
- system addresses 58

T

- test messages
 - requesting 39
 - selecting and receiving 39
- text compression 51
- traces, storage 8
- transactions organizing 49
- transmission, requesting pauses 52
- transmitting
 - large messages 25
 - messages 18

U

- undelivered message retention days 8
- updating library members 63
- user profile, archiving 59