Fixed Income Clearing Corporation

MBS EPN IMPLEMENTATION GUIDE

DTCC

Securing Today. Shaping Tomorrow.®

Table of Contents

The MBS EPN Implementation Guide (hereinafter, the "Guide") is provided as a convenience to Participants and for information only. Although FICC may make this Guide available to Participants, it shall be under no obligation to do so nor, having once or more done so, shall FICC have a continuing obligation to make available this Guide or other related information of a certain type.

FICC DOES NOT REPRESENT THE ACCURACY, ADEQUACY, TIMELINESS, COMPLETENESS, OR FITNESS FOR ANY PARTICULAR PURPOSE OF ANY INFORMATION PROVIDED TO MEMBERS IN THIS GUIDE, WHICH IS PROVIDED AS-IS. FICC SHALL NOT BE LIABLE FOR ANY LOSS RELATED TO SUCH INFORMATION (OR THE ACT OR PROCESS OF PROVIDING SUCH INFORMATION) RESULTING DIRECTLY OR INDIRECTLY FROM MISTAKES, ERRORS, OR OMISSIONS, OTHER THAN THOSE CAUSED DIRECTLY BY GROSS NEGLIGENCE OR WILLFUL MISCONDUCT ON THE PART OF FICC. FICC SHALL NOT BE LIABLE FOR: (1) ANY LOSS RESULTING DIRECTLY OR INDIRECTLY FROM INTERRUPTIONS, DELAYS, OR DEFECTS ARISING FROM OR RELATED TO PROVIDING THIS GUIDE; AND (2) ANY SPECIAL, CONSEQUENTIAL, EXEMPLARY, INCIDENTAL, OR PUNITIVE DAMAGES.

Further, the information contained in this Guide is subject to change. Participants and other authorized users of the Guide will find the most current version of the Guide on FICC's internet site, www.dtcc.com/clearing-services. FICC shall bear no responsibility for any losses associated with the failure of Participants or other authorized users to follow FICC's most current Guide.

1. Introduction

The Electronic Pool Notification (EPN) System is a real-time store and forward message switch that provides an electronic communication network for transmitting MBS pool allocation information quickly, efficiently, and reliably. EPN is operated by the Mortgage-Backed Securities Division (MBSD) of the Fixed Income Clearing Corporation (FICC), a wholly owned subsidiary of Depository Trust & Clearing Corporation.

This document provides a technical description of the EPN System, and includes all the necessary requirements and procedures to implement EPN within your firm. The document is divided into the following sections:

- **EPN Computer to Computer Interface** (CTCI) includes requirements and procedures for implementing EPN CTCI.
- **Disaster Recovery** includes requirements and procedures for implementing an EPN disaster recovery plan. This critical component will prevent service interruption in the event of an unrecoverable disaster at the primary EPN Site.
- **FTS/Batch Reporting** includes requirements and procedures for implementing Batch Reporting Services via RTTM Web Report Center and/or File Transmission Service (FTS).
- **Conformance Testing** includes requirements and procedures for testing all interfaces with EPN.

Readers are advised to become familiar with some additional documents that are referred to throughout this document:

- **EPN Guidelines and Codes Guide:** includes information regarding several key EPN functions, as well as the usage and effects of populating several key fields as well as eligible codes for DK, Cancel and Pool Substitution messages. Also included are the Acknowledgment codes issued on the Application Acknowledgment message (CTCI only).
- **MBS EPN Message Layouts:** includes message layouts for processing messages via FICC's Electronic Pool Notification (EPN) system.

2. EPN Computer-to-Computer Interface (CTCI)

2.1. General Information

EPN Computer-to-Computer Interface (CTCI) is a message exchange protocol that provides on-line, real-time, automated message exchange with EPN. CTCI employs Transmission Control Protocol/Internet Protocol (TCP/IP), which is a consistent set of rules that determine what, when, and how data may be exchanged over a communication link.

2.1.1. User Requirements and Responsibilities

CTCI Participants are responsible for developing their Application Program Interface (API) and completing any necessary changes to internal allocation and/or optimization applications. Once development and thorough in-house testing are completed, the application must be demonstrated to satisfy specified Conformance Testing requirements.

2.1.2. Codesets and Codeset Conversion

The internal codeset used by the EPN host is the American Standard Code for Information Interchange (ASCII). This codeset is also used for CTCI via TCP/IP communication with user computer systems.

It is crucial to recognize the difference between an EPN session and the physical connection of a PC and telecommunication line. There is not a one-to-one relationship between sessions and communications lines for TCP/IP users. One line may simultaneously carry one, two, three or four sessions.

Each EPN session is associated with a unique Connection ID and message routing table. To establish a session, a user must send a Logon message (LO) to EPN. The message contains a Subscriber ID, Connection ID, Account ID, Password, and Connection Type (TCP/IP Input Only = I, TCP/IP Output Only = O, TCP/IP Both Input and Output = B). EPN references the Connection ID to identify the session and the Account ID to locate the message routing table entry that applies to the session.

Except where noted (e.g. the Start and End of Message indicators), only printable alphanumeric characters should be used in the data portion of EPN messages. **However, because EPN neither edits nor validates every data field, recipients should be capable of receiving messages with embedded non-printable characters.** Messages containing non-printing characters (binary) are forwarded unmodified from EPN to the contraparty without any special indication. The message recipient is responsible for error detection and subsequent action.

EPN Messages must begin with one Start of Message character (Hex 02) and terminate with one End of Message character (Hex 03). If either character is embedded within the body of a message, the EPN Service will produce the following responses.

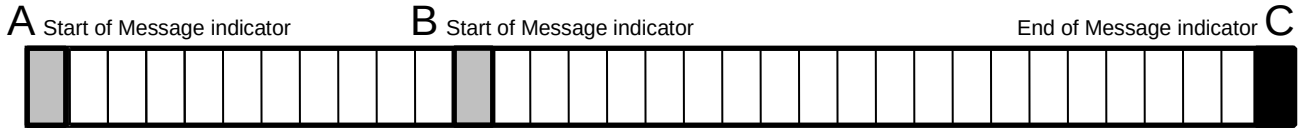

When a Start of Message indicator (B) appears after the true Start of Message (A) and before the End of Message indicator (C), the bytes from A to B are discarded, and a message is written to the internal error log. The bytes from B to C are sent to the message processor for edit checking.

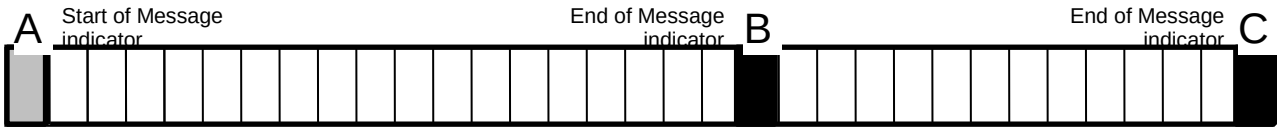

When an End of Message indicator (B) appears before the true End of Message (C), the bytes from A to B are interpreted as a complete message and sent to EPN for edit checking. The bytes from B to C are discarded and an AA is not sent. A message is written to the internal error log.

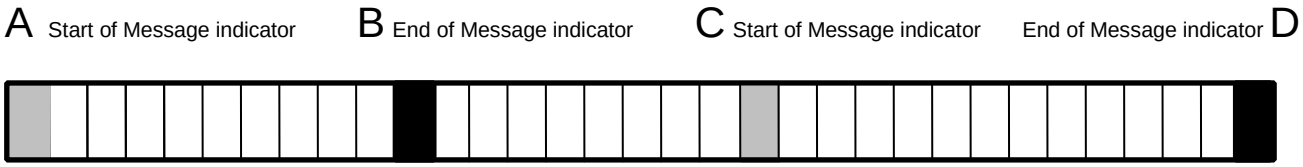

When a Start of Message indicator and End of Message indicator are missing, the bytes after the most recent End of Message indicator and before the next Start of Message indicator are discarded. In this case, the bytes delimited by B and C are discarded. A Start of Message indicator is anticipated immediately after position B.

2.1.3. Subscriber ID, Accounts, and Connections

Each EPN user has one Subscriber ID, which can be associated with multiple accounts. Accounts are used to route messages to discrete business groups (i.e. books of business). Input messages must include the contra party's Account ID (Contra).

Accounts and outbound (from EPN) connections might not have a one-to-one relationship. Messages from several inbound accounts can be directed to a single outbound queue. Conversely, each account can direct messages to different outbound queues for different message types (e.g., ONs, DKs, CC, and CXs).

Message routing is configured on a per-account basis, so any permutation is possible. For example, consider an EPN CTCI subscriber with two accounts, MBXX and MBYY. Business messages intended for account MBXX could be routed to connection 1, and non-business messages (such as AAs) could be routed to connection 2. Likewise, business messages intended for account MBYY could simultaneously be routed to connection 3 and non-business messages to connection 4.

Whether or not to utilize this flexibility is left to the subscriber's discretion. Multiple accounts can also have the same configuration, and routing would be determined purely by message type. AAs could be sent to one Connection ID (for all accounts), and ONs sent to another, and so forth.

The diagram illustrates the relationship between Subscriber ID, accounts, outbound queues, and Connection IDs.

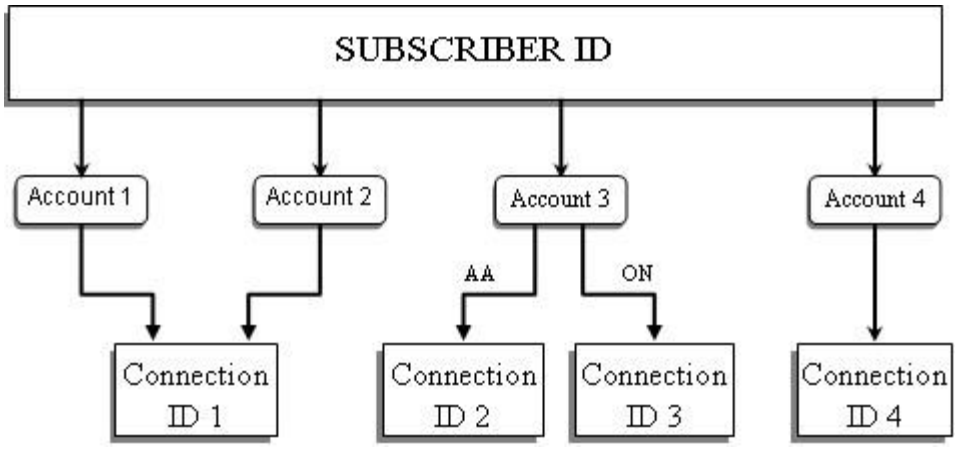

Figure 1: Subscriber ID, accounts, and Connection IDs

Connection IDs serve two functions:

- Define the account from/to which messages are sent
- Provide security

2.1.4. Message Types

CTCI employs seventeen message types, which are grouped into two categories:

- Business Messages
  - Original Notification (ON), EPN input
  - Original Notification (ON), EPN output
  - Don't Know (DK), EPN input
  - Don't Know (DK), EPN output
  - Cancel (CX), EPN input
  - Cancel (CX), EPN output
  - Pool Substitution (CC), EPN input
  - Pool Substitution (CC), EPN output
- Administrative Messages
  - Logon (LO), EPN input
  - Application Acknowledgment (AA), EPN output
  - Status Request (SR), EPN input
  - Text (TX), EPN output
  - Retransmission Request (RR), EPN input
  - Heartbeat Primary (HP), EPN input
  - Heartbeat Acknowledgment (HA), EPN output
  - Last Sequence (LS), EPN input
  - Last Sequence (LS), EPN outbound

During processing, EPN adds a Message ID, timestamps, descriptions for submitter and contra, output sequence number, and time-stamp to the ON, DK, CC, and CX input messages. These messages do not pass straight-through. For additional details, refer to MBS EPN Message Layouts.

The user CTCI application must send messages that are valid for EPN input. The conditions under which messages may be sent or received are described in the following paragraphs.

2.1.5. LO (Logon) Message (Administrative Message)

To initiate a session, the CTCI user application must send a valid Logon (LO) message, which includes the mandatory Subscriber ID, Connection ID, Account ID, Password, and Connection Type fields. These fields must contain the correct MBSD-assigned values for the associated CTCI connection. Subscriber ID is the same for all sessions established by a user, but the Connection ID and Password are session specific. The Password is associated with the Connection ID only, not with the Subscriber ID.

An LO message also includes a Connection Type field that indicates the session type: TCP/IP Input only = I, TCP/IP Output only = O, TCP/IP Both Input and Output = B. The LO message must provide the appropriate Connection Type, or data loss can occur. For example, if an LO message specifies Connection Type ***Input Only = I***, but AA messages are routed to the Connection ID, the AA messages cannot be received.

The CTCI user application must log on to EPN at the start of every session. Therefore, an LO message must be the first message sent to EPN. Any subsequent LO messages received during a session are taken and the previous connection is dropped.

EPN does not employ a logoff message, and there is no dedicated procedure to end a session. Disconnecting a TCP/IP session (sending a close command, connection time-out, or hardware failure/DR event) functions as a logoff. When a session is disconnected or lost, the user is considered logged off. To regain access, the user must log on by sending a new LO message.

EPN acknowledges receipt of an LO message by sending an AA to the appropriate Connection ID, as defined by the message routing table. Keep in mind that the source of the LO and the destination of the AA might not be the same Connection ID. For example, if an LO is sent from a Connection ID with a Connection Type ***Input Only = I***, the AA must be sent to a different Connection ID (AAs cannot be sent to Input Only Connection IDs; must be bidirectional or Output Only). Refer to Section 2.1.18 for further details.

2.1.6. ON (Original Notification) Message Input - (Business Message)

An ON message contains specific information regarding the pools being delivered to satisfy a TBA trade obligation. The message is an official notification of intention to deliver. ON messages (input) can be transmitted to EPN after CTCI user application logon.

Upon receipt of an ON message, EPN performs editing, validation, and safe storage. The message is also delivered to the designated recipient. For CTCI, each ON message (input) must contain an EPN Message Header and EPN Pool Detail (1-250 lines).

2.1.7. DK (Don't Know) Message Input - (Business Message)

A DK message can be sent in response to an ON/CC message to inform the sender that the information will not satisfy good delivery at settlement. The recipient of the DK message, or the DKed party, should always be the sender of the ON/CC message to which the DK message applies. In addition, the Target Message ID field of the DK message should contain the Message ID of the ON/CC message being DKed.

EPN does not match or cross-check DK messages against ON/CC messages. It is therefore possible to receive a DK message that does not correctly identify the associated ON/CC message. Furthermore, a DK message can be in response to a notification that was sent by phone or fax, and the Target Message ID field will be blank.

DK messages are processed and delivered in the same manner as ON messages. First, a DK message is delivered to the DKed party, then EPN also sends an acknowledgement (AA) to the DK message sender.

For CTCI, each DK message must contain an EPN Message Header and EPN Pool Detail (1-250 lines).

2.1.8. CX (Cancel) Message Input - (Business Message)

A CX message notifies the recipient of a previous ON message or a CC message (replacement pools) that the message was sent in error and is being canceled. The sender of the CX message should always be the sender of the ON message to which the CX message applies.

CX messages, like DK messages, should target the original message, but EPN does not match or crosscheck CX messages against original messages. Therefore, it is possible to receive a CX message that does not correctly identify the associated ON or a CC message. Furthermore, a CX message could be in response to a notification that was sent by phone or fax, and the Target Message ID field will be blank.

CX messages are processed and delivered in the same manner as ON messages. The CX message is delivered, then EPN also sends an acknowledgement (AA) to the sender of the CX message.

For CTCI, an input CX message must contain an EPN Message Header and EPN Pool Detail (1-250 lines).

2.1.9. CC (Pool Substitution)

The EPN Pool Substitution service message enables members to simultaneously cancel an existing pool allocation and replace it with a new one using a single EPN message. The message is an official notification of Cancel of the original allocation and delivery of the intention to deliver replacement pools. Participants are advised to refer to the EPN Guidelines and Codes document for further information on submitting CC messages.

Upon receipt of an CC message, EPN performs editing, validation, and safe storage. The message is also delivered to the designated recipient.

For CTCI, each CC message (input) must contain an EPN Message Header and EPN Pool Detail (1-250 lines).

2.1.10. AA (Application Acknowledgment) Message - (Administrative Message)

For each LO, ON, DK, CC, and CX message received, EPN returns an application acknowledgment (AA) to the sender. The specific message that has been acknowledged is identified by the input sequence number (Original Sequence Number) and user identification (Internal ID) included in the AA. These numbers allow the AA recipient to identify the message for which the AA applies.

AAs can be positive or negative. Identified by an ACK Code field entry of 0200, a positive AA indicates that an acknowledged message has been successfully edited/validated, safestored, and logically delivered to the recipient. In other words, the acknowledged message can be accessed by the AA recipient. A user implementation must have the capability to recognize and record successful transmission of positively acknowledged messages.

Negative AA acknowledgments are sent in response to invalid messages. A negative AA is identified by an ACK Code value other than 0200, which also identifies the first invalid field data. User implementations must have the capability to recognize, report, repair (if possible), and retransmit negatively acknowledged messages. Refer to the EPN Guidelines and Codes Guide, for specific code definitions.

AAs are routed to the session that is designated in the user message routing table as the AA destination. The destination session of the AA can be different than the source session of the message that was acknowledged. All AAs for an account are sent to the same Connection ID.

2.1.11. SR (Status Request) Message Input - (Administrative Message)

An SR message requests the status of unacknowledged messages. Although unlikely, a communication, hardware, or system outage can result in a lost or undeliverable AA acknowledgement for a specific message. If a method to detect such conditions is not provided, the user implementation would be unable to determine the status of a message previously input to EPN.

All user implementations must be equipped with programmable logic to detect messages that remain unacknowledged for a designated period, or time-out value. Time-out value must be user-programmable between 1-128 seconds, with a default of 60 seconds. If one or more messages remain unacknowledged after the time-out value has expired, the user implementation must automatically generate an SR message that can query the status of up to five messages.

EPN responds to a valid SR with a discrete AA acknowledgement for each message defined in the SR status request. Each AA, though delayed, is the same message that had been returned in response to the original message(s). EPN does not send an AA for the SR message itself. A positive AA (ACK Code 0200) indicates that an original message was accepted, and a negative AA indicates that the original message was rejected or not found.

When a negative AA is returned, the original message must be retransmitted by the sender. The same is true for messages that fail input edit/validation, are corrected and resubmitted, then fail again. These are the only acceptable conditions for retransmitting input messages. All retransmitted messages must contain a new sequence number.

If EPN does not respond to an SR after an additional time-out period has expired, the user implementation should send an additional SR (recommended maximum of 3). The user implementation should alert operating personnel to contact MBSD after unacknowledged SR attempts.

An SR message can originate from the same session as the original message(s) or from a different session. User-side CTCIs must populate the Original Connection ID field on all SR messages.

2.1.12. RR (Retransmission Request) Message - (Administrative Message)

An RR message is a user-side CTCI output that automatically requests retransmission (current day only) of an EPN outbound message that was not received or retained by the user.

EPN responds to a valid RR message by retransmitting the requested message(s). EPN does not send a positive AA acknowledgment for a valid RR message, but a negative AA acknowledgment is generated for an invalid RR request.

A retransmission request message must include the Connection ID and the outbound (from EPN) sequence number. For example, if a message with outbound sequence number 123 was sent from EPN to Connection ID 2110, but the message was lost, the user-side CTCI can send a retransmission request. The RR must specify Connection ID 2110 and the appropriate sequence range (e.g., from outbound sequence number 123 to outbound sequence number 123). Note that the RR message can originate from a Connection ID different from the Connection ID where the lost message was sent (2110). In other words, the lost message was sent to Connection ID 2110, but the RR can originate from a different Connection ID.

Performance considerations dictate that a maximum range of five messages can be requested in a single RR. Although five is the current maximum, this value is subject to change. MBSD will notify all EPN users if the maximum range is scheduled for modification. Sequence range must be specified as an inclusive value. For example, an RR for sequence numbers 123-127 is valid, but an RR for sequence numbers 123-128 exceeds the maximum of five, triggering EPN to generate a negative AA (reject code 0223).

A negative AA response to an RR message indicates a possible problem and should be investigated. As a general rule of thumb, an RR message should not be sent to EPN until a retransmitted message or negative AA has been received for each outstanding RR message. This practice is referred to as RR Pacing. It is also good practice to take advantage of sequence number ranges. A single RR message with a sequence range of five is processed faster and more efficiently than five individual RR messages.

IMPORTANT NOTE: Retransmission is not limited to business messages. Administrative messages, such as old negative AAs, can be retransmitted if the corresponding Connection IDs and sequence numbers match. To distinguish a negative AA that has been successfully retransmitted and a negative AA that was sent in response to an invalid RR that is requesting retransmission, compare the outbound sequence number and the original outbound sequence number.

2.1.13. TX (Text) Message - (Administrative Message)

Text (TX) messages notify users of a gap in the input sequence numbers. User implementations should report the contents of the Message Text field to operators and/or administrators. TX messages are also sent during a disaster recovery scenario. See Section 3, Disaster Recovery, for further details.

2.1.14. HP (Heartbeat) Messages - (Administrative Messages)

During idle periods, when the user does not generate message flow, Heartbeats are used to verify connectivity between a user implementation and the EPN application. The EPN Alarm Processor must detect a Heartbeat or any message within a specified timeframe, or the inactivity may be interpreted as a connectivity problem. User implementations must be equipped with logic to automatically generate heartbeat messages at a specified interval. The interval period must be user-programmable from 1-60 minutes (minimum), with a default setting of 10 minutes.

2.1.15. Heartbeat Message Sequence Numbering

As an additional safeguard, EPN monitors the sequence numbers of incoming HP messages along with other messages. If the sequence deviates from an incremental succession, a TX message is sent to alert the user of a possible error. The anticipated sequence number is incremented upon message processing.

2.1.16. HA (Heartbeat Acknowledgment) Message - (Administrative Message)

For each Heartbeat message it receives, EPN returns a Heartbeat Acknowledgment (HA). The HA includes the sequence number of the message being acknowledged. A flag is used to declare that the system is operating normally or that a disaster has occurred and the Alternate Site is receiving data. See section 3, Disaster Recovery for additional information. Sequence numbers are assigned, in succession, to all HA messages, along with all other messages, that are sent in response to HP messages. Heartbeat Acknowledgements are sent to the session defined in the message routing table.

2.1.17. LS (Last Sequence) Message - (Administrative Message)

An LS message requests the sequence numbers of the last sent and last received messages between EPN and a specific Connection ID.

If a connection is lost, message sequencing must be determined. The user implementation can send an LS message to EPN and query the sequence numbers of the last messages processed for a specific Connection ID. Provided that an LS message is valid, EPN returns an LS reply, which includes the sequence state of the specified Connection ID prior to LS message processing.

Each LS reply is assigned the same sequence number as its corresponding LS message. Replies are not incrementally sequenced and therefore gap detection is not performed. Sequencing continues for subsequent messages as if the LS message and corresponding reply did not occur.

When an invalid LS message is received, EPN returns a negative AA message. Subsequent outgoing messages are sequenced as if the negative AA message did not occur.

LS message replies are sent to the Connection ID defined in the user's message routing table as being the recipient of LS messages.

Unsolicited LS messages are also sent during a disaster recovery scenario. See Section 3, Disaster Recovery for further details.

2.1.18. General Message Flow

CTCI is a bidirectional, synchronous protocol, so each CTCI session supports simultaneous message transfer to and from EPN (receive and transmit simultaneously). There are no requirements to alternate the direction of transmission for CTCI sessions.

A specific CTCI Connection ID can be configured as Input Only, Output Only, or Bidirectional. For example, users can opt in their EPN Configuration Plan to configure a specific CTCI Connection ID for Input Only. The Connection ID would have the capability of sending messages to EPN, but it will not receive any messages (including ONs and AAs). Keep in mind that at least one active CTCI session must be maintained as Bidirectional or Output Only so that messages can be received from EPN during the operating day.

The Decoupled Acknowledgment feature of a user CTCI implementation can be configured to control message flow in one of two ways. New messages can be sent to EPN in succession without requiring any acknowledgements, or a new message can be sent only after receiving an acknowledgement for the previous message.

The following diagrams illustrate bidirectional and unidirectional message exchange. Diagram [A] illustrates one CTCI session, configured for Bidirectional message flow. The session can send and receive messages simultaneously. Diagram [B] illustrates one session configured as Input Only and one session configured as Output Only.

Bidirectional Message Flow

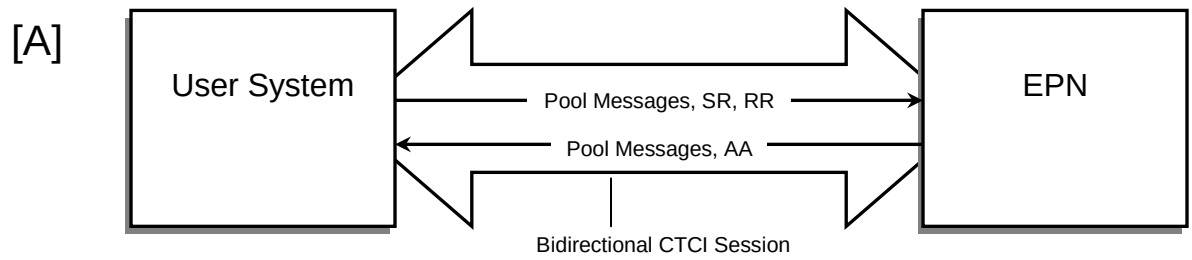

Unidirectional Message Flow

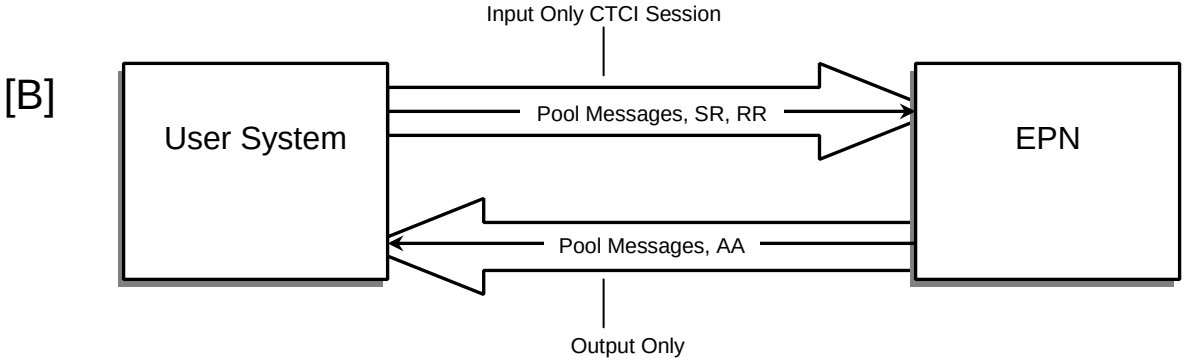

Figure 2: Bidirectional and Unidirectional Message Flow

2.1.19. Message Routing Tables

Contra is a mandatory field and must be populated for each input business message. Although this field provides the contra's EPN account number, it does not specify a particular terminal or CTCI session. Because a contra account can have several lines, each running up to four simultaneous CTCI sessions, the EPN System must have additional instructions for routing messages to a final destination. These instructions are defined in the Message Routing Table.

At the start of each CTCI session, EPN gathers information from the user's Logon (LO) message, including a Connection ID, which serves as an identifier for the session. The LO message is also used to associate the user with a specific Message Routing Table. The table contains the routing destinations, by message type, for all users within an account, as specified in the EPN Configuration Plan. Each account has a custom Message Routing Table. For each account, all AAs must be sent to one Connection ID.

2.1.20. EPN Routing Table Worksheet

The EPN Configuration Plan includes a table where implementers must define the message routing schemes that will satisfy their firm’s requirements. Several configurations are possible. The following examples provide just a few possibilities for EPN message routing for illustrative purposes only.

In the first example, Firm A elects to receive all ON, CC, and AA messages for an account on CTCI Session # 1. Firm A chooses to route CX and DK messages to CTCI Session #2. The routing table for this account configuration is shown in Table 1.

Output Message Types (from EPN)	Route To
ON	CTCI Session #1
CC	CTCI Session #1
DK CX	CTCI Session #2
AA	CTCI Session #1

Table 1: Message Routing for Firm A

Firm B elects to receive ON and AA messages for an account on a CTCI Session #1. DK, CC, and CX messages will be routed to CTCI Session #2. The routing table for this account configuration is shown in Table 2.

Output Message Types (from EPN)	Route To
ON	CTCI Session #1
DK CX CC	CTCI Session #2
AA	CTCI Session #1

Table 2: Message Routing for Firm B

In the last example, Firm C elects to have all business messages and AAs for an account routed to a single CTCI session. The routing table for this account configuration is shown in Table 3.

Output Message Types (from EPN)	Route To
ON CC DK CX AA	CTCI Session #1

Table 3: Message Routing for Firm C

2.1.21. Message Sequence Numbering and Gap Detection

Introduction

CTCI protocol incorporates sequence numbering for input (to EPN) and output (from EPN) messages, and sequence number gap detection is performed for input messages by Connection ID. Any in-transit message loss is detected, and the message sender is notified (either EPN or the user).

All input and output CTCI messages must be sequentially numbered in the Sequence Number field. A separate sequence is maintained for input messages and output messages.

At the start of the operating day, the sequence number for EPN outbound messages and the expected sequence number for input messages are reset to 1. In other words, at the start of the operating day, EPN anticipates that the first message it receives from every user CTCI session will have a sequence

number of 1. Therefore, the sequence number for a user's first message sent to EPN, which is a Logon message (LO), must have a sequence number of 1 for all CTCI sessions or the user will receive a gap detection TX message.

Sequence numbers are assigned by the message sender, whether EPN or the user, and each Connection ID has its own sequence. If there is a second logon for the same Connection ID — for example, after a session is lost — the sequence numbering resumes from the previous session. The sequence is not reset from one session to another for the same Connection ID during the same business day.

If, for example, a firm has multiple Connection IDs, each with a separate numbering sequence, there must be additional identifiers to distinguish message 1 for Connection ID-One and message 1 for Connection ID-Two. Therefore, every message output by EPN is identified by a combination of Subscriber ID, Connection ID, Internal ID, and sequence number.

The maximum count for a numbering sequence is 999,999. If more than 999,999 messages are sent to or from a Connection ID during the EPN operating day, the sender count must roll over to sequence number 1. However, it is highly unlikely that a rollover would ever occur because industry volumes do not remotely warrant such numbers.

If EPN detects a gap in input sequence numbers, it sends a Text (TX) message to alert the sender, as defined by the message routing table. Because of the potential financial liability involved in exchanging allocation information, EPN processes out-of-sequence messages rather than rejecting them. Upon receipt of a Text (TX) message, the user CTCI application must alert its operations personnel and take action (if possible) to resolve the problem. However, **if EPN detects a sequential duplicate sequence number, the message identified by the duplicate number is rejected,** and a negative application acknowledgement is sent (negative acknowledgement is sent regardless of the Possible Duplicate Flag being set).

Input Sequence Numbering

EPN internal processing monitors sequence numbers based on a compare-and-increment algorithm, using an internal counter (Expected Sequence Number) for each Connection ID. When an input message is received, the received sequence number and the expected sequence number are compared. If the received sequence number does not match the expected or the previous sequence number (a sequential duplicate), EPN sets the expected sequence number to the received sequence number plus 1. The received message is then processed, and a gap detection TX message is returned to the sender with the following details.

EXPECTED SEQ # XXXXXX, RECEIVED SEQ # YYYYYY (blanks fill the field to 130 characters)

However, if a business message is received with a sequence number that matches the last received message, the message is rejected. Only sequential business messages with duplicate sequence numbers are rejected. For example, if four business messages with sequence numbers 2, 3, 2, and 3 are received, all of the messages are processed. However, if four business messages, all with sequence number 5 are received (i.e. sequence number 5, 5, 5, 5), the last three business messages are rejected.

It is critical for user CTCI applications to maintain correct input sequence numbering and to send messages in sequence number order. Otherwise, sequence numbering and gap detection are ineffective and message transfer integrity is lost. But even with gap detection, it is possible that messages can be lost in-transit, without any indication from the communications protocol. Furthermore, if duplicate input sequence numbers go undetected (i.e. the previous example of 2, 3, 2, 3), the corresponding AAs for those messages are ambiguous, and the sending user system may be unable to determine which AAs apply to which messages.

Cross-Referencing Input Sequence Numbers: Business, AA and SR Messages

Each AA message includes the sequence number and any Internal IDs for the acknowledged message. Users can cross-reference AAs to original messages using the original (input) sequence number, Internal ID, or both.

An SR message must contain the sequence number(s) of the message(s) to which the query applies because SR messages can apply to a range of messages, and Internal IDs are not required to be in sequence (as sequence numbers are). In fact, Internal IDs are not a mandatory field, and could be blank.

Output Sequence Numbering

Output sequence numbering allows for gap detection at the user system if any messages are lost in transit. To ensure correct sequencing, it is important for user CTCI applications receiving messages to provide a compare-and-increment algorithm similar to EPN.

As with user-generated input sequence numbers, EPN generates output sequence numbers on a Connection ID basis. All EPN output messages for a specific Connection ID are sequence numbered--including TX, AA, and output business messages.

Possible Duplicate Messages

At the beginning of each operating day, sequence numbering, which is generated by the sender, must start at number 1. From that point on, the incremental succession of a sequence must not be reset or altered during the operating day. EPN will always follow this model, even after a service interruption. In which case, if the transmission status is uncertain, a message may be retransmitted with the Pos Dup field set to X (normally blank), indicating that it may be a duplicate message. This should only occur in the event of a regional disaster and the system recovers at the Remote Site.

When a user implementation receives a possible duplicate message, it must check EPN's Message ID field. If the field content is identical to the Message ID of a previously received message, the user implementation should recognize and discard the message as a possible duplicate. If the message ID field is unique, the user implementation should process the message as normal.

Similarly, if the transmission status of an input message is uncertain, a user can resend the message with the Pos Dup field set to X, alerting the contra-side of a possible duplicate.

2.1.22. Security

To log on to EPN, users must have correct combinations of user identifiers and passwords, which are provided upon approval of the EPN Configuration Plan.

2.2. EPN Internal Processing

This section of the Implementation Guide outlines EPN internal message processing, and will prepare implementers for the task of designing and implementing a CTCI application in their firm. This section outlines ON, DK, CC, and CX message processing.

T1 is applied when the complete message is received for processing.

T2 - After the message has passed all edits and is processed by the Message Processor, the T2 stamp is applied. **The T2 time-stamp is used to determine good delivery**.

T3 is applied to the message when it is sent to the outbound table by the Message Processor.

T4 represents the time the message is stored internally by the Message Processor.

T5 represents the time a message is written to the outbound table by the Output Formatter.

T6 represents the time the message is successfully delivered a message over the communications line.

2.2.1. Selecting Input Messages for Processing

EPN member firms can have multiple accounts, each simultaneously running up to four CTCI sessions, and all contending for message processing. Therefore, a fairness methodology is employed to ensure that all firms receive equal treatment. The Subscriber Selection Process processes one message from each subscriber in a sequential fashion according to an internal subscriber list.

The first step in message processing is to determine the next message to be read. The Subscriber Selection Process loops through a stored list of current EPN subscribers, and once the next submitter is identified, that submitter's oldest message is passed to a message processor. The message is then marked as having been passed to a processor. The Subscriber Selection Process starts at approximately 7:00 am and remains active throughout the EPN processing cycle.

2.3. CTCI via TCP/IP

2.3.1. TCP/IP Fundamentals

General Information

Connections are initiated by the user CTCI implementations. The EPN System is the host/server for TCP/IP connections, and line handlers adhere to TCP protocol standards. Line handlers are dual purpose, receiving user initiated messages and sending outbound messages.

Physical Layer Parameters

TCP/IP based CTCI communication is established via SMART, Radianz, and SFTI Networks. The EPN host is equipped with DSUs/modems connected to network routers, and the communication lines are configured for full-duplex, synchronous, or asynchronous serial transmission.

For information on the SMART, Radianz, and SFTI Networks, please contact Integration at ficcintegration@dtcc.com

---

Connection and Message Exchange Procedures

TCP/IP communication in CTCI is only via Berkeley Socket connections and procedures. The Berkeley socket type for TCP/IP is designated Sock-Streams. The logical port number for connecting to EPN via TCP/IP is 2001. Telnet, UDP, and FTP are not utilized in CTCI.

The procedure for opening a TCP/IP session, exchanging data, and closing the session using the Berkeley sockets API are outlined below. Implementers can use Sockets, Streams, WinSock, or any other applicable API.

1. The EPN host issues a *socket ( )* call to acquire a socket endpoint.
2. EPN prepares to receive incoming connection requests using the *bind ( ), listen ( ), and the blocking accept ( )* calls.
3. The user issues a *connect ( )* call to port 2001, which generates an open request.
4. EPN receives the open request, and the EPN program returns from the *blocking accept ( )* call.
5. EPN waits for data from the user's program.
6. EPN listens for additional incoming requests while processing the data associated with the last request.
7. EPN blocks on a *recv ( )* call waiting for transactions.
8. The first transaction should be a Logon message.
9. Data is exchanged, via *send ( )* and *recv ( )*, in the direction specified by the Connection Type.
10. To end the session, the user CTCI implementation issues a *close ( )* on each socket that was opened for CTCI transmissions. This halts all processing.
11. EPN responds by issuing a *close ( )*.

Because TCP/IP provides data exchange via continuous streams of data rather than discrete blocks or records of data, blocking/deblocking procedures are not used. Messages are delimited within the stream by a Start Message Indicator (hex 02) and an End Message Indicator (hex 03).

Table 4 summarizes the calls issued in the Connection and Message Exchange Procedures.

User application	**EPN Server**
Initialization	initialization
open ( )	open ( )
	bind ( )
	listen ( )
connect ( )	accept ( )
send ( ) / recv ( )	recv ( ) / send ( )
recv ( ) / send ( )	send ( ) / recv ( )
shutdown ( ) (optional)	shutdown ( ) (optional)
close ( )	close ( )
Cleanup	cleanup

Table 4: TCP/IP Utilization

TCP does not support half-duplex communications, but a CTCI session can be configured as unidirectional (simplexed) by restricting specific connection IDs to receive or send only (see Section, 2.1.19, Message Routing Table). Send-only or receive-only sessions follow the identical process as described above.

2.4. EPN CTCI User Application Implementation Guidelines

2.4.1. EPN CTCI User Application should include the following capabilities:

1. Send LO message first and wait for AA message sent by EPN server
   Once the CTCI user application is started it should try to establish a connection with EPN Server by sending LO (Logon) message to EPN server. After the LO message is sent, the CTCI user application should wait for AA (Application Acknowledgement) message from EPN Server. The CTCI user application will not process any further messages before receiving an AA message corresponding to the sent LO message. Please note that AA message will be sent to a different connection ID if LO message was sent from Connection ID that is not set up to receive AA messages. Therefore, the EPN CTCI user application should be able to look for LO messages at all Connection IDs assigned to the Account(s) the CTCI user application is implemented for.
   Should EPN CTCI user application try to send any business messages before establishing connection with EPN server (i.e. before receiving acknowledgement of receipt of LO message sent), the messages should be “held” by EPN CTCI user application and sent to EPN server only after receiving AA message from EPN server confirming that LO message was received and a session with EPN server has been established.
   If AA message is not received after an agreed interval, the EPN CTCI user application should send a new LO message. The CTCI user application should repeat the process till AA message is received.

2. Receive messages from EPN server
   Once the AA message is received confirming that connection was established, the EPN CTCI user application is ready to send and/or receive messages using the current Connection ID. In addition to sending messages, the user application should be able to receive responses sent by EPN server. The responses to business messages might be received by another Connection ID, depending on Account setup.

3. Respond to connection issues.
   If connection between EPN CTCI user application and EPN server is lost during EPN CTCI user application execution, the user application should be able to detect the interruption. Once the interruption is discovered, EPN CTCI user application will send a LO message to re-establish connection and proceed as outlined above (1. Send LO message first and wait for AA message sent by EPN server)

4. Send a HP message at a specified interval.
   EPN CTCI user application should have logic to constantly send a HP (Heartbeat) message to EPN server so connectivity between the user application and EPN server is verified. The interval period the message is sent should be between 1 and 60 minutes, with 10 minutes being a default. Once the HP message is received by EPN server, the server will send HA (Heartbeat Acknowledgement) message to EPN CTCI user application.

2.4.2. EPN CTCI user Application Best Practices:

Recovery-Restart Considerations

There are significant differences between the EPN compare-and-increment algorithm and the algorithm that users must implement. These differences are most notable if EPN experiences a service interruption and must perform a recovery-restart procedure.

The first case where this applies is when a message is sent from EPN, but for some reason the message is not safestored. For example, this can occur when an EPN Message Processor delivers a message to an outbound connection, then malfunctions before the message is safestored. In this case, EPN would not output an AA, so the original message sender must send an SR to determine message status. At that point, EPN outputs a negative AA (message not received), and the sender must retransmit the message. Now, a duplicate of the original message has been created.

Because EPN is unaware of the duplicate, it passes the retransmitted message to the recipient without a Pos Dup indication. Therefore, preventive measures must be taken whenever an EPN service interruption initiates a recovery-restart sequence. In these instances, senders must populate the Pos Dup field when retransmitting a potentially lost message, and manually reconcile with the recipients.

Another potential problem that can occur after an EPN restart-recovery is that a user may receive messages with duplicate output sequence numbers that are not flagged as possible duplicates. This can occur when EPN is able to safestore a message, but the final time-stamp is lost during a service interruption. The user application must check these messages and take the appropriate action, processing a good message or discarding an already received message. See Disaster Recovery Implementation for further restart-recovery information.

3. Disaster Recovery (DR)

FICC provides cost-effective services for its member organizations. An integral component of these services is adequate contingency planning, evidenced by the redundant system backups that safeguard EPN processing.

All EPN users, including those with a CTCI implementation are required to maintain at least one complete backup configuration to access EPN.

3.1. Recovery Design

The Fixed Income Clearing Corporation maintains three fully functional data centers, located in different geographical regions:

- Primary Site
- Alternate Site
- Remote Site

Each site is equipped with a single highly available virtual application/server running on a redundant hardware cluster. Should one of the hardware hosts fail, the virtual application/server gets moved to one of the remaining hardware hosts and the application resumes processing. This transfer is referred to as intra-site vMotion failover.

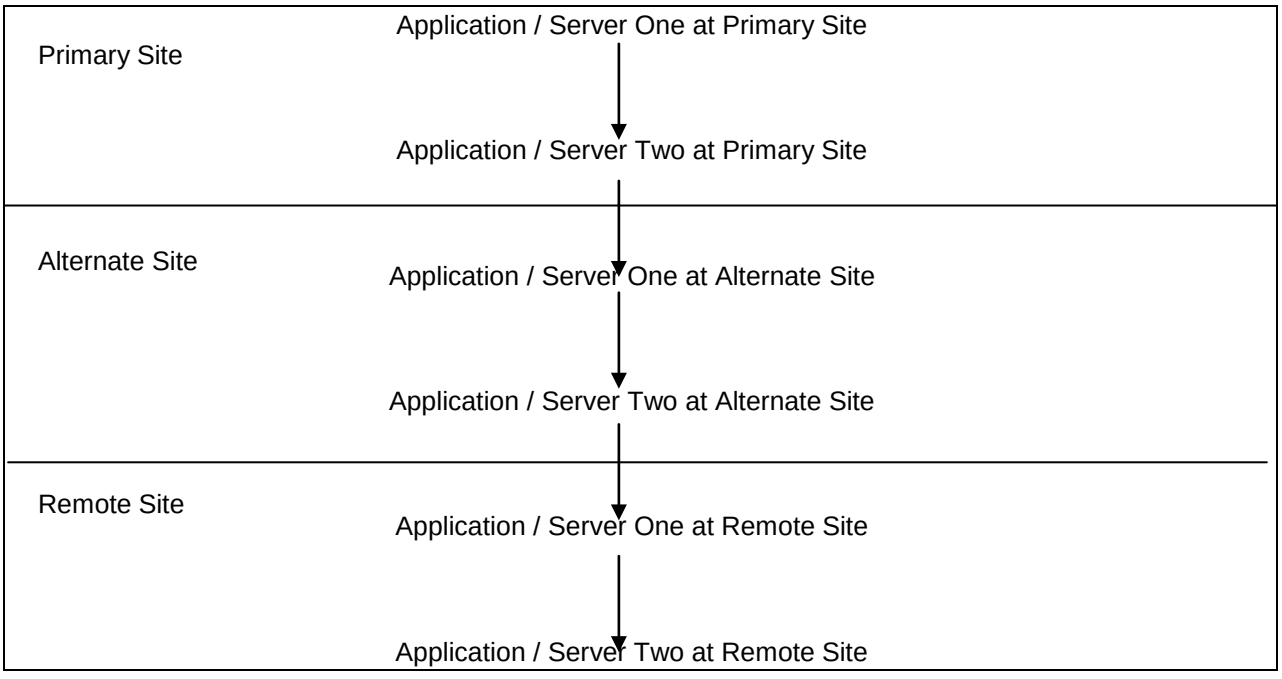

Intra-site failover will be transparent to the end user and connectivity will not be impacted. During an inter-site failover, all connectivity is lost, and participants must log in to re-establish a communication link. All three Sites utilize the same IP address and ports. Therefore, participants only need to retain a single IP address for the Primary, Alternate and Remote Site.

Traffic will failover to the Remote Site only if a disaster encompassing both the Primary and Alternate Site has occurred and it has been determined that all MBSD applications as well as all other DTCC applications must operate remotely. As mentioned above, the IP address for the Remote Site is the same as the Primary and Alternate Site. For information, please contact Participant Interface Planning (PIP) at 888-382-2721 (select Option 5, followed by Option 5).

3.2. Failover Notification

Intra-site failover at the Primary or Alternate Site is essentially transparent to the user, and participants will not receive formal notification. As previously noted, connections to EPN are lost during inter-site failover. Consequently, the EPN system will not transmit acknowledgements for new status requests or outstanding messages. Should this occur, participants must send an LO message to reestablish a connection.

If EPN responds to an LO message with a positive acknowledgement, failover to the application/server at the Alternate Site has been successful, and the participant can resume message transmission. Failure to receive a positive acknowledgement indicates that failover is still in progress. In this case, participants must continue to send LO messages until a positive acknowledgement is received.

An intra-site failover does not initiate a formal notification. Participants should not anticipate an incoming LS message, the HA flag is not set to 2, and EPN does not send a text message. If a gap detection message is received upon re-establishing a connection to the system, the participant should send an RR (or LS) message to identify missing messages (i.e., messages that were sent by a participant but not received by EPN).

3.3. Disaster Recovered

A disaster is considered recovered when the Alternate Site is Primary-ready. Upon re-establishing a connection to EPN, users are notified, via messages, that the system has been recovered in DR mode. First, heartbeat acknowledgment (HA) messages, which are now sent from the Alternate Site, have a Disaster Level of 2 (Disaster Level field = 2). And second, a Disaster Recovered TX (Text) message and an LS (Last Sequence number) message are sent to all Subscribers, but only to the connection IDs that have been configured to receive these message types. If a single connection ID is configured to receive all TX and LS messages for an entire firm (all accounts), users must be notified of the Disaster Recovery status.

The Disaster Recovered Text message that is broadcast upon recovery includes the following:

Text Type: 04 *Disaster is recovered and EPN is running at the Disaster Recovery Site.*

The LS messages that are automatically generated upon disaster recovery provide the last sequence number for each connection ID. Sequence numbers corresponding to the Alternate environment should be used for reconciliation purposes as described below. Although LS messages are automatically generated upon disaster recovery, user-initiated LS requests are still processed as normal.

Each user application should compare the sequence number reported in the LS message with the expected value, revealing any sequence gaps, or potentially lost messages. If warranted, business messages should be re-sent to EPN with the Possible Duplicate Message (Pos Dup) indicator activated (set to X). For example, if a user application database indicates that the last inbound sequence number is 104, but the last inbound sequence number reported in the LS message is 100, business messages 101104 should be re-sent with the Pos Dup indicator activated.

Messages that are re-sent in the Alternate Site do not require the same inbound sequence numbers that were originally sent to the Primary Site. The next inbound sequence number should be determined by the LS message for that connection ID.

Each user application should also compare the last outbound sequence numbers reported in the LS messages with the expected values. Based on this information, each user should submit Retransmission Request (RR) messages for the missing messages.

A discrepancy in outbound sequence numbers does not necessarily indicate that business messages have been lost. Outbound sequence numbers are incremented for each Acknowledgment message sent from the Primary Site. If no business message was sent out on the corresponding connection ID at the

Primary Site before a disaster is declared, the non-business traffic will cause an outbound sequence number discrepancy in the Alternate region.

3.3.1. Disaster Recovered: Miscellaneous

In addition to setting the disaster level in HA messages to 2 and sending TX and LS messages, the Message ID field is incremented by 100,000. This field (format: mmddxxxxxxx) uniquely identifies any message on the system (where mm = current month, dd = current day, xxxxxxx = incrementing number for every business message processed). For example, if the last business message that is known to have been processed at the Primary Site is assigned message ID 06250001234, the next message processed, after recovery, at the Alternate Site is assigned 06250101235. This ensures the process of uniquely identifying messages that may have been re-sent as part of the disaster recovery procedure.

The User Test environment is available during normal processing days, and continues to be available even in the event of a declared disaster.

After a disaster has been recovered, the Alternate Site will continue to process Primary traffic for at least the remainder of the day. Should the Primary Site become available, it will not be utilized until the next day.

3.3.2. Alternate Site Available

Primary Site, Inactive	**Alternate Site, Active**
inactive	Traffic will be directed to the Alternate Site for active processing of Primary messages.
Inactive	Business messages are sent and received
Inactive	HP messages are received and HA messages are sent to users
Inactive	Disaster Recovery flag on HA is set to two (Alternate Site available)
Inactive	Message ID incremented by 100,000
inactive	TX message with text type ‘04’ is sent indicating disaster recovered. LS messages sent indicating last sequence numbers.

3.3.3. Miscellaneous Alternate Site

On the day of a site failover, several indicators are triggered: EPN sends an LS message and a text message, and the HA flag is set to 2. However, EPN may continue to operate from the Alternate Site for several days. On subsequent days, the HA flag remains set at 2, but no other messages or indicators are sent. The first EPN message ID is reset to mmdd0000001*,* and sequence numbering is as usual (i.e., outbound sequence number for each connection ID starts a one).

3.4. Failover to the Remote Site

As previously noted, traffic will failover to the Remote Site only if a disaster encompassing both the Primary and Alternate Site has occurred and it has been determined that all MBSD applications as well as all other DTCC applications must operate remotely. This being the case, the failover sequence will be Primary Site to Alternate Site to Remote Site, as outlined in Section 3.1. Once a failover to the Alternate Site has occurred and a disaster has been declared, including all related notification, an additional failover to the Remote Site will occur. Should EPN failover to the Remote Site after a disaster has been declared from the Alternate Site, no additional disaster declaration or associated messaging will be sent. The HA

flag will be set to 0 while EPN operates from the Remote Site. Recovery of the Primary Site will reset the HA flag to 0, indicating that EPN is operating normally.

Note that in the event regional disaster, the system may not recover at the Alternate site but rather, it may recover directly at the Remote Site. If this occurs, then the notification events that occur during a recovery at the Alternate site will not be triggered.

4. FTS/Batch Reporting

4.1. Introduction

EPN batch reporting is available to users via the existing MBSD batch reporting mechanisms. Users who are not receiving reports from the MBSD Trade Processing System must establish a new telecommunications link (requires hardware and software installation). It is important to note that EPN Service and the MBSD Trade Processing System employ different hardware and software configurations.

FTS generates data as machine readable output (MRO) in bulk format, which is not possible through interactive processing. MRO is easily integrated into member firms' internal data processing applications.

IMPORTANT NOTE: For general information about FTS/Batch Reporting or to set up participant access, please contact the Integration team at ficcintegration@dtcc.com . Please allow at least two weeks for setup processing.

The following EPN output files are available through FTS:

13150381MBSCC COB EMESDETC

13150382MBSCC COB EMESSUMC

13150384MBSCC EOD EMESDETE

13150385MBSCC EOD EMESSUME

The EPN Detail file (13150381MBSCC COB EMESDETC, 13150384MBSCC EOD EMESDETE) provides a replay of all messages delivered to or from the account on the day the file was generated. The Close-of-Business file details message until 15:00:00, and the End-of-Day file details messages delivered for the entire day.

The EPN Summary file (13150382MBSCC COB EMESSUMC, 13150385MBSCC EOD EMESSUME) provides a summary of all messages by TBA CUSIP and message type delivered to or from the account on the day the file was generated. The Close-of-Business file details message until 15:00:00, and the End-of-Day file details messages delivered for the entire day.

5. Conformance Testing

5.1. General Information

To ensure a successful implementation, users must have competency in executing EPN functions. MBSD requires all users to demonstrate competency by executing the applicable Conformance Test Packets prior to EPN participation. Conformance Test packets are available from MBSD for all EPN technical interfaces:

- EPN Computer-to-Computer Interface (CTCI)
- Report Center Connectivity
- EPN/MBSD File Transmission (applies only to participants that are both Clearing and
- EPN members)

5.1.1. Testing Phases

Each Conformance Test consists of three phases:

- **Connectivity Test:** Verifies line connectivity by executing handshake, loop back, and protocol testing. Configuration, or *Handshake*, testing ensures that a working communication link exists between the CTCI user system and EPN.

  The required data communication equipment must be installed prior to handshake testing. Users are responsible for all cable connections between their data communication equipment and their computer systems.
- **Message Integrity:** Scripted test cases demonstrate a user's ability to perform various EPN functions. Data is used to generate test Close-of-Business and End-of-Day Reports available via FTS and MBSD's on-line report inquiry feature.
- **Internal Application Integrity:** Users submit copies of Primary messages and receive all EPN output, which provides an opportunity to identify any internal processing deficiencies in a risk-free testing environment.

5.1.2. Schedule

Conformance Testing must be completed prior to using EPN in the Primary environment. Typically, MBSD will schedule Conformance Testing to coincide with the installation of the communication lines. MBSD must be notified, in advance, if additional time is required to prepare for internal testing. All communication lines, hardware, software, and pre-tested internal applications must be in place prior to scheduling Conformance Testing. Dedicated resources must be committed for the duration of the test.

5.1.3. Connectivity Testing Procedures

DTCC Participant Interface Planning (PIP) technicians will assist participants in configuring the various parameters and options associated with data communication devices. For CTCI circuits, PIP technicians test the communication link, analyzing the binary data transfer rate from EPN to the user and back. This is referred to as Bit Error Rate (BER) testing.

During loop back and BER testing, PIP technicians monitor data communication/data terminal (DCE/DTE) interfaces at the EPN host site, checking for the correct timing of various handshake signals.

Once these tests are successfully completed, the assigned EPN and MBSD host port is configured for operation. As a final test, the user must successfully establish and maintain a session with EPN for CTCI, as well as EPN and MBSD for terminal and printer services.

5.1.4. Message Integrity Testing

Upon completion of EPN Connectivity Testing, users must demonstrate their capability to send and receive each message type. MBSD has created test scripts to simulate EPN activity within a controlled input/output environment. Creating these scenarios in your test environment will ensure the integrity of your application and exercise EPN functionality. This aspect of testing requires a significant effort, and several hours (minimum) of hands-on activity should be anticipated.

Do not confuse this phase of testing with volume or performance testing. For overall control, verification, and ease of troubleshooting, the amount of test data is limited. EPN will not accept input from users who exceed the data requirements for this phase of testing.

Reconciliation: User staff must be dedicated to performing prompt and accurate reconciliation of input/output.

Scripted Data: MBSD requires all firms to execute scripted tasks. Users assume the role of buyer and seller while submitting and receiving messages, both valid and invalid. Once a testing session has been established, the user will transmit the specified messages. EPN will accept the messages and transmit a response. Users must verify the results immediately.

5.1.5. Internal Application Integrity Testing

After completing the MBSD scripted tests, users can also generate non-scripted messages in the test environment. MBSD encourages users to send all message types and exercise all functionality (e.g., send a message with 250 Pool Detail lines).

5.1.6. Conformance Testing

Connectivity	**CTCI**	**Web Access**
1. Configuration of devices	✓	
2. Loop back Testing	✓	
3. BER Testing	✓	
4. Monitor Signals	✓	
5. Host Port Configuration	✓	✓
6. Establish/Maintain Session	✓	✓

Message Integrity - Scripted Data	**CTCI**	**FICC Report Center**
1. Logon	✓	✓
2. Send ON, DK, CC, CX	✓	

3. Send SR, RR	✓	
4. Receive ON, DK, CC, CX	✓	
5. Receive TX, AA	✓	
6. Access all screens		✓

MESSAGE INTEGRITY- Non-Scripted Data	**CTCI**	**FICC Report Center**	**MBSD FTS**
1. Logon	✓	✓	
2. Send ON, DK, CC, CX	✓		
3. Send SR, RR	✓		
4. Receive ON, DK, CC, CX	✓		
5. Receive TX, AA	✓		
6. Print EOD/COB reports		✓	
7. Receive EOD/COB Files			✓