

Fault Tolerance

This topic describes iLink session fault tolerance.



Customers must wait for in-flight resend requests to be fulfilled before logging out or failing over to the backup gateway.

Fault tolerance in a network environment is characterized by rapid recovery from such failures such as process termination, hardware failure, or network disconnects.

One method of providing fault-tolerance is through a mechanism called Failover; the goal is to minimize service interruption caused by error conditions.

Process failure may be caused by hardware failure of the machine on which the process runs, as well as software error; faulty logic, and improper memory handling, among others. Network failure includes loss of TCP/IP connectivity between the client application and the iLink FIX Gateway. This may be caused by such items as faulty network interface cards, frayed wires, and power failure on network components (i.e. routers).

The iLink FIX Gateway initiates a controlled failover when it detects either process or network failure that impacts its ability to service the client.

See the following topics for details regarding each Fault Tolerance error condition and example scenarios:

- [Summary of Fault Tolerance Requirements on Client Applications](#)
- [Clients Choosing Not to Implement Fault Tolerance](#)
- [Fault Tolerance Implementation](#)
 - [Logon Procedure with Fault Tolerance](#)
 - [Examples of Fault Tolerance Scenarios](#)
 - [Client System Sends FTI Status of 'U' for Beginning of Week or Mid-Week Logon](#)
 - [Application Message Sent Over a Backup Connection](#)
 - [Backup Client System Sends Incorrect FTI](#)
 - [Client System FTI Status Assigned as Primary or Backup](#)
 - [Client System Process Complies with FTI Instruction](#)
 - [iLink 2 Assigns FTI Status](#)
 - [Primary Client System Sends Incorrect FTI](#)
- [Fault Tolerance Error Conditions](#)
 - [Client Primary Process Failure](#)
 - [Client Backup Process Failure](#)
 - [Primary CGW Failure](#)
 - [Backup CGW Failure](#)
 - [Primary MSGW Primary Failure](#)
 - [Backup MSGW Failure](#)
 - [Network Failure](#)
- [Duplicate Session Connection Attempts](#)

Summary of Fault Tolerance Requirements on Client Applications

The guidelines for implementing fault tolerant client applications are:

- Coordinate applications such that the primary and backup processes each establish a separate and independent content stream to the iLink Gateways via TCP/IP socket connection.
- For the [Beginning of Week](#) and [Mid-Week Logon](#), set the Fault Tolerance Indicator (FTI) in tag 49-SenderCompID field to 'U' for undefined.
- For the In-Session Logon, set the FTI in tag 49-SenderCompID field to 'P'.
- Examine the FTI in tag 56-TargetCompID for each incoming message. Depending upon its value, the client application must behave accordingly.
- Send application messages (e.g. New Order - Single, Order Cancel/Replace Request) only through the primary content stream where sequencing is enforced per FIX 4.2 protocol.
- Send only administrative messages over a backup content stream. Communication over a backup content stream is for link maintenance via administrative messages only. Application messages sent over a backup connection are rejected.
- Messages exchanged over a backup content stream serve to maintain only the connection. As a result, sequence numbers are not relevant and are "0" (zero).
- Maintain tracking of inbound and outbound sequence numbers between the primary and backup client processes to ensure proper failover processing.

Clients Choosing Not to Implement Fault Tolerance

iLink 2 does not mandate that client systems use the fault tolerance feature; however, CME Group strongly recommends using this functionality. Clients who do not implement the fault tolerance cannot dynamically recover from process or network failures.

For customers who do not want to implement fault tolerance, CME Group provides an optional method for using the fault-tolerance indicator. Customers not implementing fault-tolerance on their applications may submit their logon request with an FTI of 'N'. The CME Group responds with a logon confirmation message with an FTI of 'N' and expects all subsequent messages during that session to contain an FTI of 'N'. If the customer submits a message with an FTI of anything other than 'N', the customer is logged out.

This functionality also prohibits a customer from attempting to log in with another application. If a customer requests multiple iLink Gateways for fault-tolerance availability, but chooses to submit a logon message with an FTI of 'N', any attempts to authenticate through another iLink Gateways result in a logout of those requests from the FIX Gateway. Client systems will receive a Logout (tag 35-MsgType=5) message with tag 58=**Primary session fault-tolerance indicator = 'N'. Backup session not allowed. Logout forced.** The existing, authenticated session is not affected.

Fault Tolerance Implementation

Clients who choose to use fault tolerance must coordinate their application processes for establishing separate and independent content streams on iLink via TCP/IP socket connections. In a typical deployment scenario, multiple redundant processes are spawned from the same executable file and each of those processes runs on a separate machine.

Important

CME Group does not recommend running redundant processes on the same machine because if a machine fails, all the processes running on it fail simultaneously.

iLink 2 has a designated host that is primary and another that is designated backup. Customers must successfully log on to the primary before attempting to log on to the backup. If a customer logs in to the backup gateway and is not already logged into the primary gateway, client systems will receive a Logout (tag 35-MsgType=5) message with tag 58=**Invalid logon. Must be logged on to Primary. Logout forced.**

Customers can only trade through primary connections. Customers can connect to the backup, but can send only [Administrative](#) messages (e.g., Test Request, Heartbeat). [Application](#) messages sent to the backup are rejected.

Disconnecting or logging out from the primary gateway—either client- or exchange-initiated—will trigger an automatic logout or disconnect from the backup gateway as well, except in the Primary gateway failure scenario.

Logon Procedure with Fault Tolerance

When the client sends a [Logon](#) message, the Fault Tolerance Indicator (FTI) in tag 49-SenderCompID must be set to 'U' for undefined. Tag 49-SenderCompID and tag 56-TargetCompID are 7 characters long and are composed of 3 sub-fields:

- **Session ID** (left-most 3 characters)
- **Firm ID** (next 3 characters)
- **Fault Tolerance Indicator** (right-most character).

Beginning of Week Logon and Mid-Week Logon (tag 35-MsgType=A) messages must be sent with the FTI in tag 49-SenderCompID set to 'U'. If the client application submits a Logon (tag 35-MsgType=A) message and the FTI is not set to 'U', a Logout (tag 35-MsgType=5) message is issued and the connection is dropped. Because In-Session Logon messages may be sent only on the primary channel, the FTI must be set to 'P'.

 All client applications, both primary and backup members, must examine the FTI in tag 56-TargetCompID for each incoming message.

Based on the value of the FTI contained in tag 56-TargetCompID, the client application must populate the FTI in tag 49-SenderCompID with the same value for all outgoing messages.

1. If the FTI is set to 'P', then the application must behave as the active member representing the fault tolerant group.
2. If the client application submits a message on the primary connection with an FTI value of 'B' in tag 49-SenderCompID, iLink 2 sends a Logout (tag 35-MsgType=5) message on the Primary channel.
3. If the FTI is set to 'B', then the application must behave as a backup member.
4. If the client application submits a message on the backup connection with an FTI value of 'P' in tag 49-SenderCompID, iLink 2 sends a Logout (tag 35-MsgType=5) message on the backup channel.
5. If a client application submits a message **without** either a 'P' or 'B', it receives a Logout (tag 35-MsgType=5) message.

The client application must acknowledge that it has successfully received and processed the FTI instruction from iLink 2 by sending the FTI in tag 49-SenderCompID for each message to CME Globex.

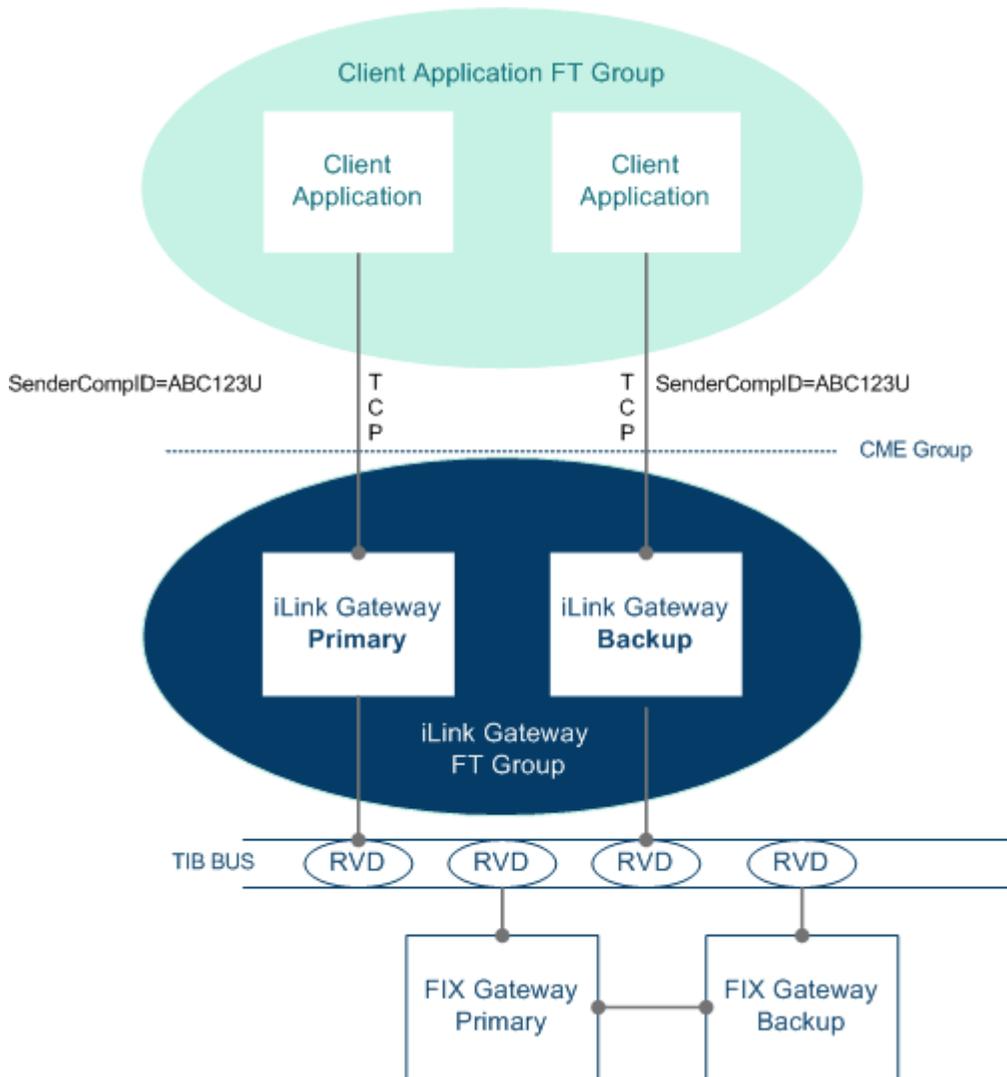
Application messages (e.g., New Order - Single, Order Cancel/Replace Request) must be sent only through the primary content stream where sequencing is enforced per FIX 4.2 protocol.

Communication over the backup is solely for link maintenance. Only administrative messages (Logon, Logout, Heartbeat and Test Request) are sent through the backup. Sequencing on the backup is not enforced; message sequence numbers in the administrative messages are zero.

Examples of Fault Tolerance Scenarios

Client System Sends FTI Status of 'U' for Beginning of Week or Mid-Week Logon

The following diagram illustrates how member processes of a client application fault-tolerant group connect to CME Globex. In this example, both client member processes send Logon messages with the FTI set to 'U' in tag 49-SenderCompID.

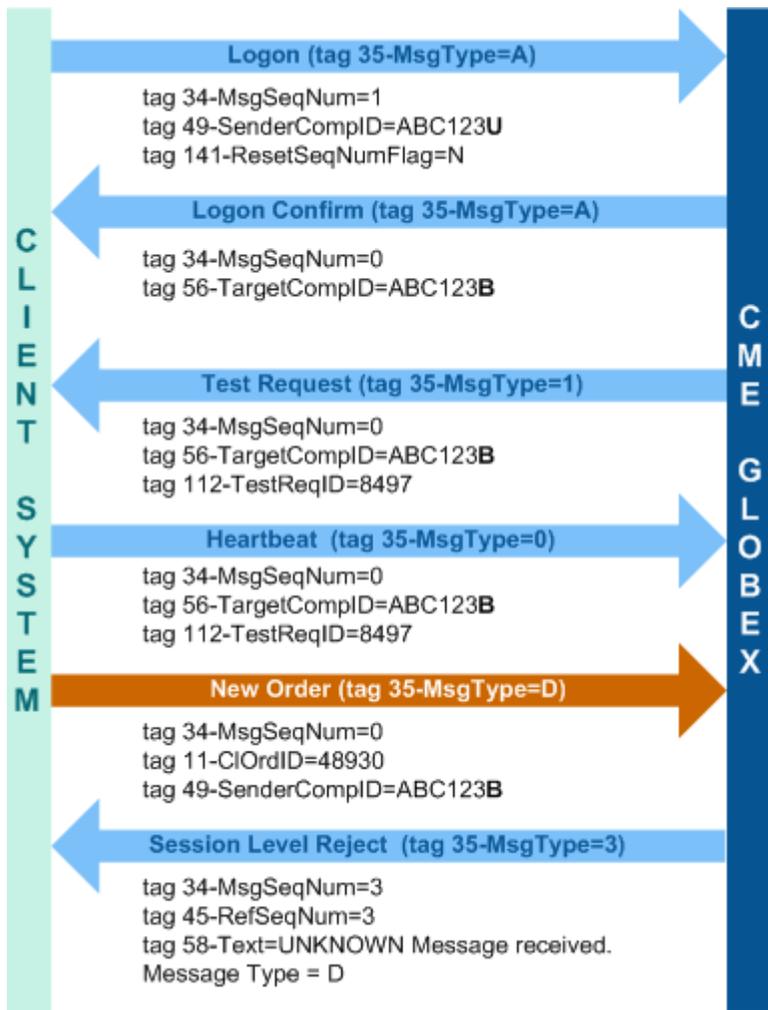


[top](#)

Application Message Sent Over a Backup Connection

In the following diagram:

- A client system is successfully logged on.
- iLink 2 elects this application process as a backup.
- The backup client system sends a [iLink 2 New Order](#) message over the backup connection (Communication over the backup content stream is for iLink maintenance only via administrative messages; application messages such as New Order – Single are not allowed).
- As a result, iLink 2 sends the backup client application a [Session Level Reject \(tag 35-MsgType=3\)](#) message with tag 58-Text containing "UNKNOWN Message received. Message Type = D".

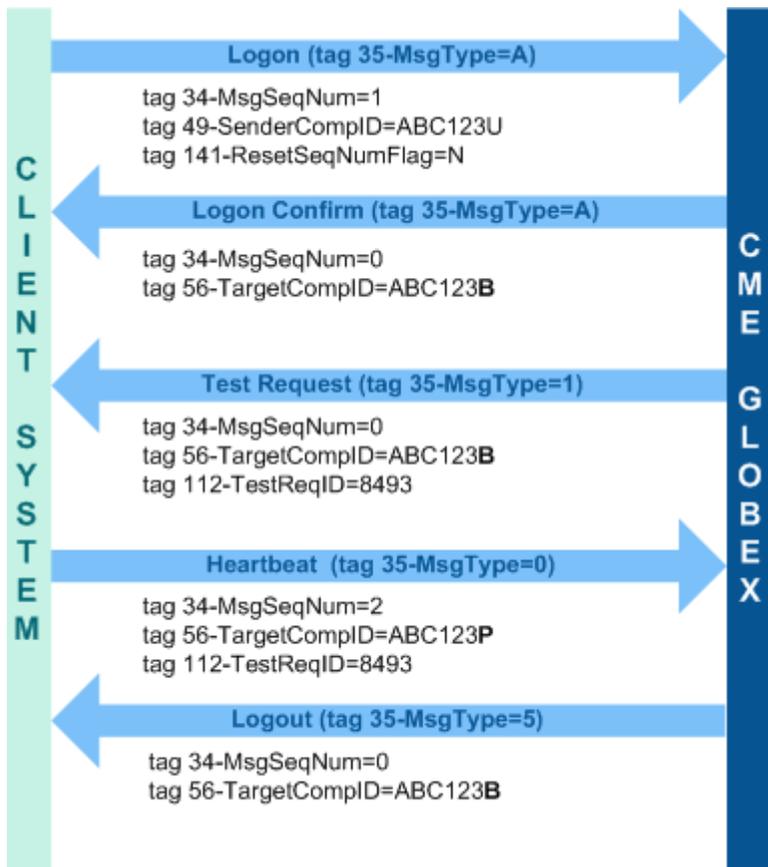


[top](#)

Backup Client System Sends Incorrect FTI

In the following diagram, the client application is logged on successfully and is designated as a backup by iLink 2:

- The backup client system sends a [iLink 2 Heartbeat](#) message with an incorrect FTI in the tag 49-SenderCompID. (The backup client system sets its FTI to 'P' instead of 'B'.)
- As a result, iLink 2 logs the backup client system out. The status of the primary client system connectivity remains unchanged.



[top](#)

Client System FTI Status Assigned as Primary or Backup

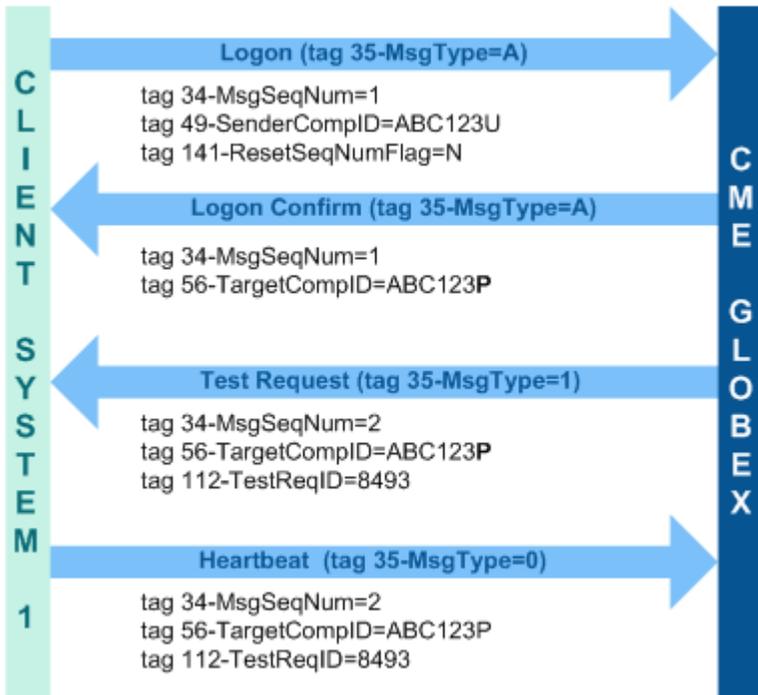
The following diagram illustrates the message flow for a successful [Beginning of Week Logon](#) scenario for two client applications side-by-side:

- Each client application examines the FTI within the TargetCompID of the Logon Confirmation message.
- The client system in column 1, on the left receives an FTI of 'P' and becomes a primary member.
- Any subsequent message sent by the primary client application must set the FTI in the SenderCompID to 'P'.
- The client system in column 2, on the right receives an FTI of 'B' and becomes a backup member.
- Any subsequent message sent by the backup client application must set the FTI in the SenderCompID to 'B'.
- This example also applies to [Mid-Week Logon](#).



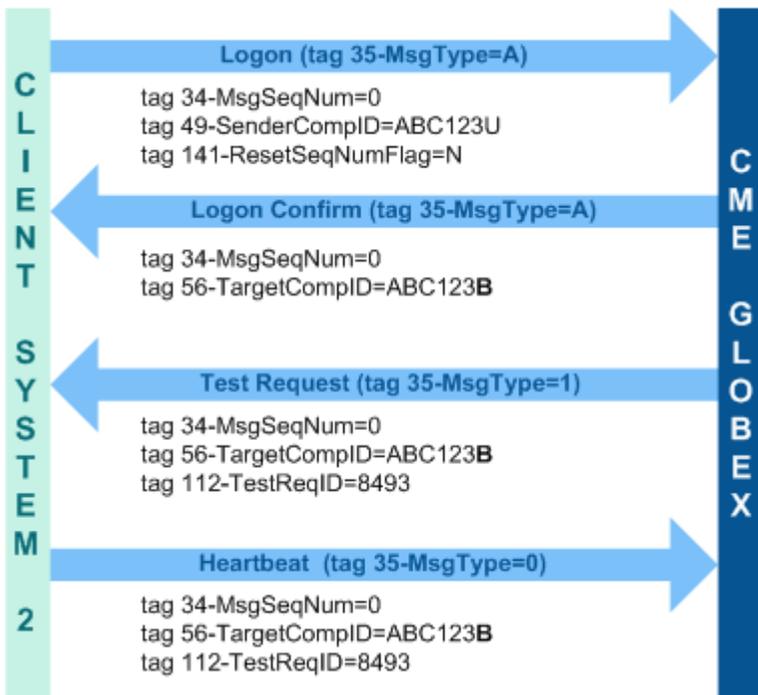
Tag 34-MsgSeqNum sent by the backup member always = '0'.

The following message scenario shows the Client System 1 as Primary.



[top](#)

This messaging scenario shows Client System 2 as Backup.

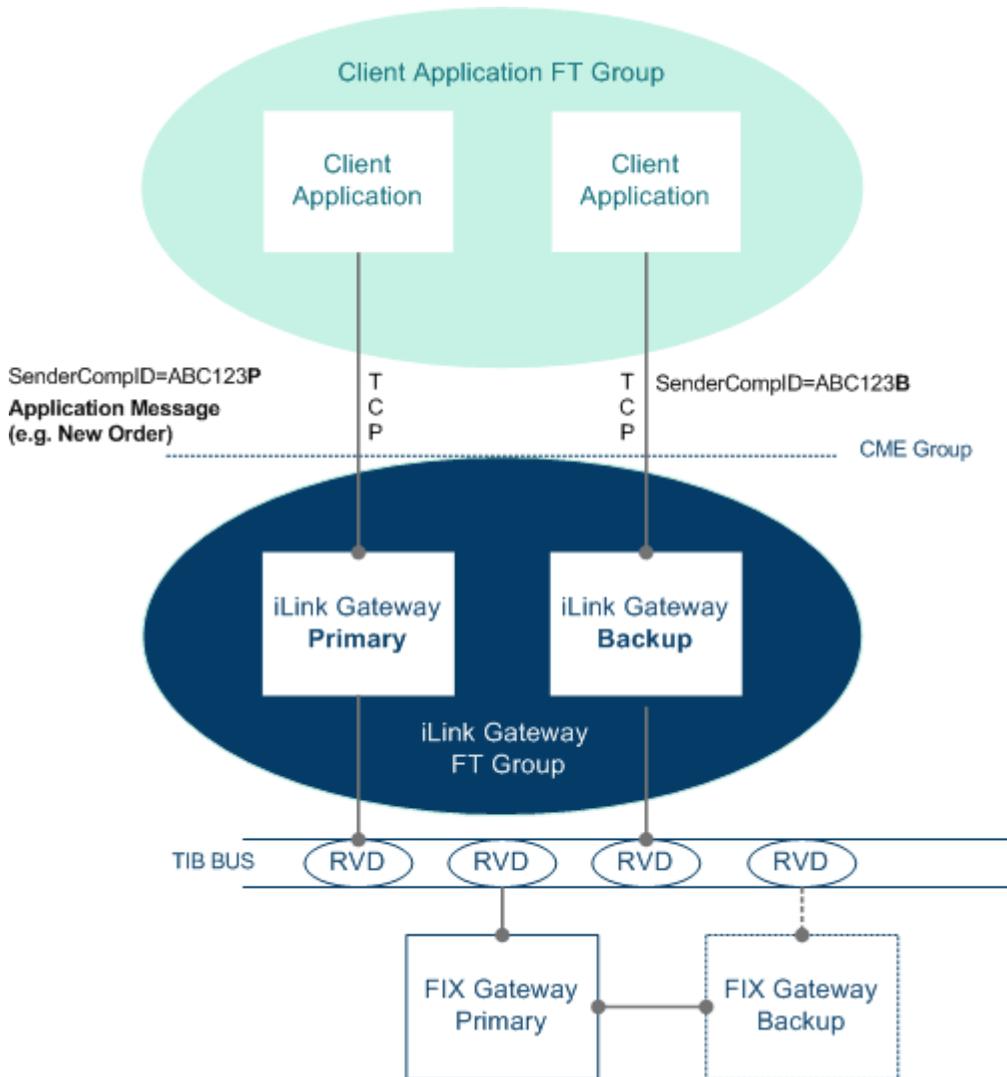


[top](#)

Client System Process Complies with FTI Instruction

In the following diagram, a client application acknowledges that it has successfully processed the FTI instruction by populating the FTI in the SenderCompID for each outgoing message:

- The primary client application submits a New Order message over the primary content stream. This New Order message contains an FTI set to 'P'.
- The member process designated as the primary sends each outbound message with its FTI set to 'P' and the backup member sends each outbound message with its FTI set to 'B'.

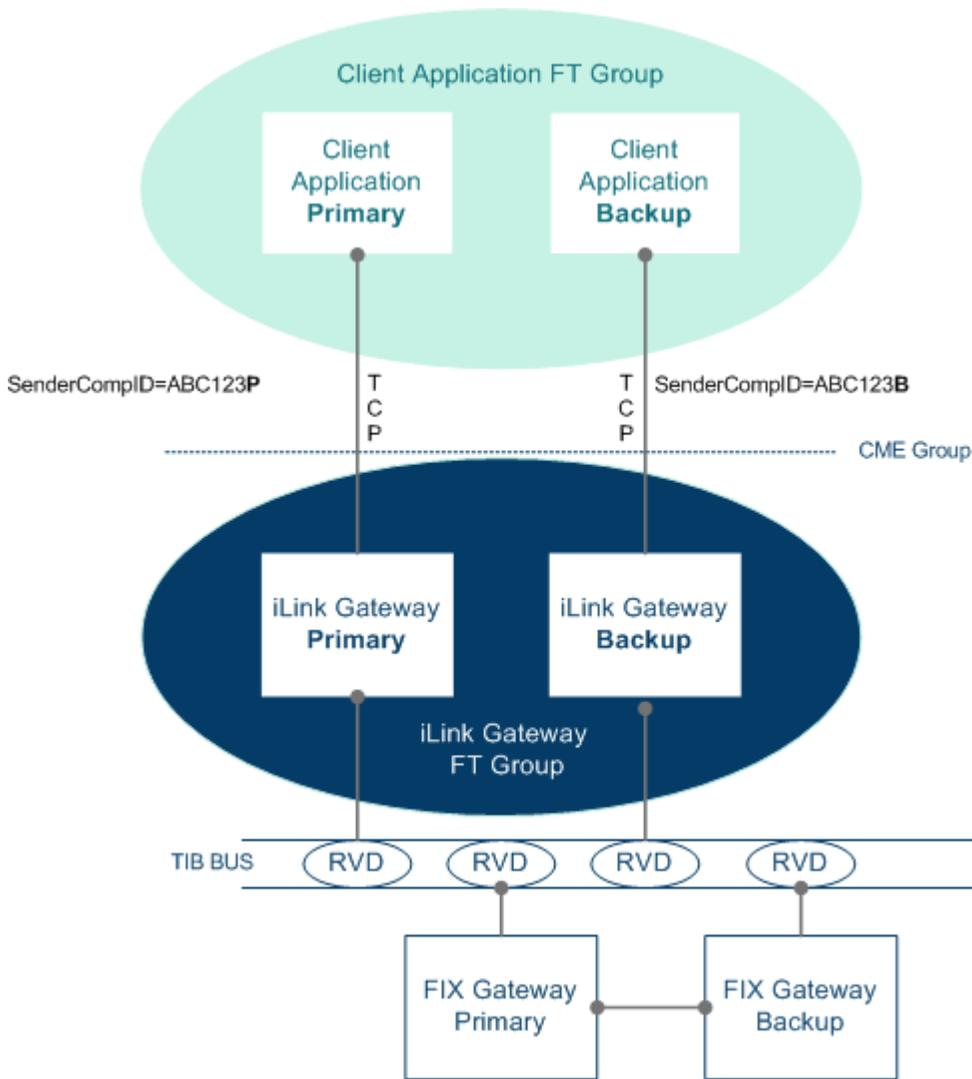


[top](#)

iLink 2 Assigns FTI Status

The following diagram illustrates how iLink 2 assigns fault tolerance status:

- As a client application is authenticated, iLink 2 dynamically assigns the fault tolerance status and populates the FTI with a 'P' or 'B' in the TargetCompID of the Logon Confirmation message.
- As all the client member processes receive and process the FTI, the fault tolerance status of the client application fault tolerant group members is fully determined.

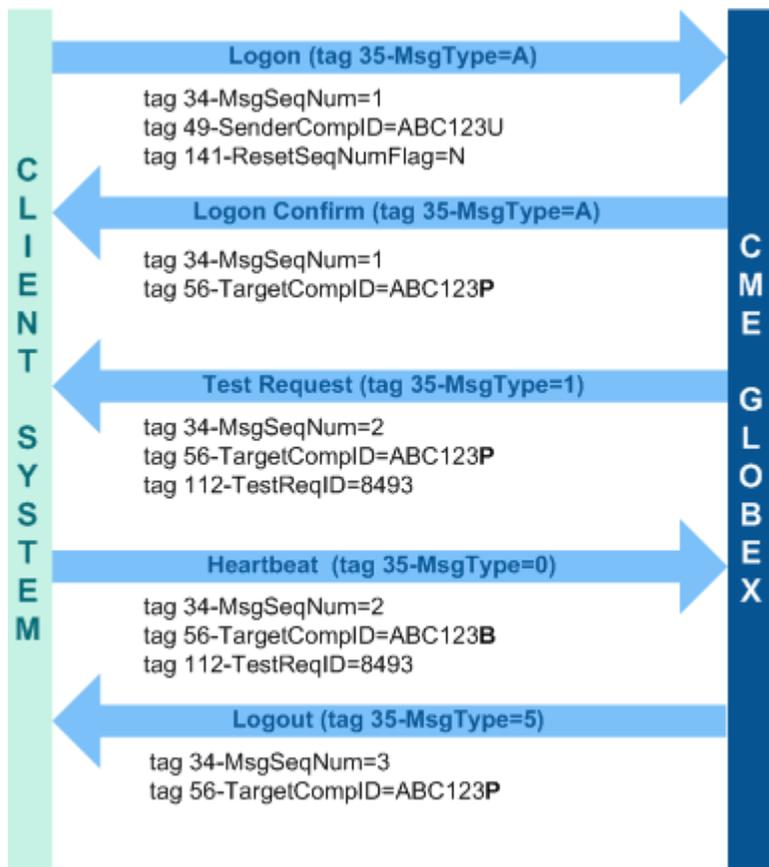


[top](#)

Primary Client System Sends Incorrect FTI

In the following diagram, the client system is logged on successfully and is designated as the primary by iLink 2:

- The primary client system sends a [iLink 2 Heartbeat](#) message with an incorrect FTI in the tag 49-SenderCompID. It sets its FTI to 'B' instead of 'P'.
- As a result, iLink 2 logs the client system out.
- When the primary client system is logged out, all the backup systems are also logged out.



[top](#)

Fault Tolerance Error Conditions

iLink 2 detects seven categories of error conditions described as follows.

Client Primary Process Failure

If iLink 2 does not receive any messages from the primary client process within the defined heartbeat interval:

- CME Globex sends a [iLink 2 Test Request](#) message to invoke a [iLink 2 Heartbeat](#) message from the client.
- If the primary client process does not respond with a Heartbeat message to the Test Request message within the defined heartbeat interval (or if the client does not send any message during the entire interval), iLink 2 designates the primary client process as failed and initiates failover.
- The primary client application is disconnected from iLink 2.
- One of the backup client applications is chosen to communicate over a new primary channel.
- The backup client application is notified of such fault tolerance status change by examining the FTI in the TargetCompID of the next incoming message.
- If the primary client process fails without closing the TCP connection, then it takes two Heartbeat intervals for iLink 2 to detect the primary process failure. The backup client application should check the FTI on every message to determine its status. If clients want to avoid the time delay in this process, then they should ensure that the TCP connection is closed whenever their application fails.

Warning

Before failover, the backup client application was receiving sequence numbers set to zero. During and after the failover process, the backup client application is responsible for ensuring that its inbound and outbound sequence numbers are synchronized with the primary application that just failed. This is critical since the newly elected primary member must know exactly where the failed member left off. If the sequence number of the message sent by the new primary client application is lower than that of the original client application, iLink 2 logs the client application out per the FIX 4.2 protocol.



If the client system is ungracefully disconnected from the primary gateway and was connected to the backup gateway, CME Globex sends a Logout (tag 35-MsgType=5) message with tag 58=**Logout initiated due to disconnect on primary connection** and disconnects the backup client system from the iLink Market Segment Gateway.

Client Backup Process Failure

If iLink 2 does not receive any message from a backup client application within a defined interval:

- CME Group sends a to invoke a Heartbeat message from the client.
- If there is no response to the Test Request message within the defined heartbeat interval (or if the client does not send any administrative message during the entire interval), the backup client application is disconnected from iLink 2.
- The status of the primary client application connectivity remains intact.

Primary CGW Failure

If the primary iLink FIX Gateway fails:

- CME Globex initiates failover by electing the ranking inactive iLink FIX Gateway to assume the primary role.
- The client application that is connected to this newly chosen iLink FIX Gateway must act as the primary for the client application FT Group.
- The client application is notified to become primary by examining the FTI in the Tag 56-TargetCompID of next incoming message.

Backup CGW Failure

iLink 2 maintains a predefined number of processes running for each iLink 2 component. If a backup iLink Gateway fails:

- The failed iLink Gateway is restored.
- The backup client application must reestablish the connection to this restored backup iLink Gateway to maintain a pool of backup content streams.
- The status of the primary client application connectivity remains intact.

Primary MSGW Primary Failure

[Click here](#) for more information on MSGW primary gateway failure.

Backup MSGW Failure

[Click here](#) for more information on MSGW backup failure.

If the backup MSGW fails:

- The failure of a backup component will not affect the primary connection state.
- The failed backup MSGW is restored.

Network Failure

In the event of network failure, iLink 2 handles socket exceptions that are thrown for network error conditions (i.e., loss of TCP/IP connectivity between the client application and the iLink Gateway). When this happens, iLink 2 designates the primary content stream as failed and initiates the failover.

[top](#)

Duplicate Session Connection Attempts

Client system cannot make a duplicate TCP connection to iLink.