

iLink Architecture

For [iLink order entry](#) on the CME Globex platform, CME provides two gateway access models.

- Convenience Gateway (CGW) - provides a connection to any CGW session, which manages routing for all order entry messages to all market segments; routing is transparent to customer.
- Market Segment Gateway (MSGW) - provides a connection to a dedicated market segment, which is defined as a group of products on a given trading engine instance.



Connecting directly to MSGWs involves a separate Session ID model, messaging behavior, and functional differences from that of the CGW.

Customers can choose to access CME Globex markets through the MSGWs or CGWs, which route through the MSGWs.

Contents

- [iLink Architecture Video Overview](#)
- [iLink Session Model](#)
- [iLink Fault Tolerance](#)
 - [iLink Fault Tolerance Scenarios](#)
- [Convenience Gateway](#)
 - [CGW FIFO Message Processing Model](#)
 - [CGW Sequence Number Management](#)
- [Market Segment Gateway](#)
 - [SFTP - MSGW Configuration](#)
 - [MSGW Sequence Number Management](#)
 - [Tag 57-TargetSubID](#)
 - [MSGW Failover Scenarios](#)
- [CME Globex Timestamps](#)
 - [Nanosecond Timestamp \(tag 5979-RequestTime\)](#)
- [Testing and Certification](#)
- [New Release Weekly Failover Testing Windows](#)

iLink Architecture Video Overview

Watch this brief video for an overview of iLink architecture.



iLink Session Model

For CGW, Session ID can exist on only one CGW.

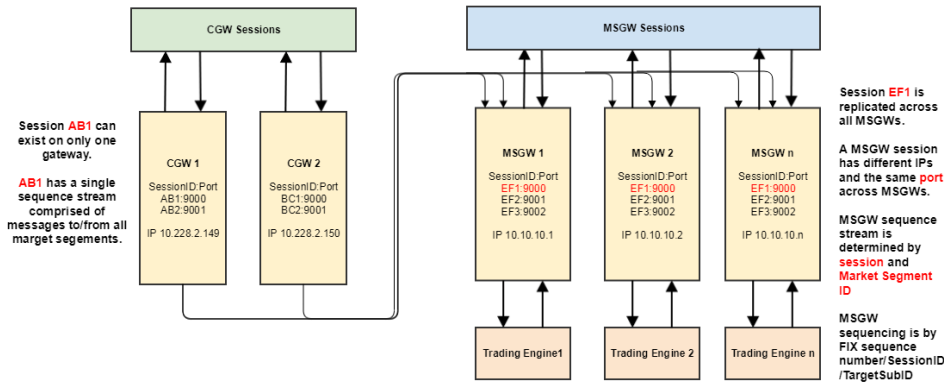
For MSGW, Session IDs are shared across multiple MSGWs. Customers utilize a single Session ID, password, and port across all MSGWs.

- Customers use their given Session ID, password, and port to connect to different MSGWs by their IP address.

- Each iLink MSGW has a unique IP address for each of the primary and backup MSGWs.

i Customers must manage the relationship between iLink sessions and corresponding IP addresses and market segments within their order management systems.

The diagram below shows how Session IDs are shared across multiple MSGWs.



! Sequence streams are unique per Market Segment ID and Session ID combination. For example, if session ABC has logged into Market Segment ID 50 and 60, then session ABC could receive Execution Reports from both market segments with the same message sequence number (tag 34-MsgSeqNum).

iLink Fault Tolerance

For iLink, Fault Tolerance is managed at the gateway application level. iLink 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 session logs on to the backup gateway before logging on to the primary gateway, CME Globex will send a [Logout message \(tag 35-MsgType=5\)](#), with tag 58-Text = **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 a backup gateway 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.
 - A Logout message (tag 35-MsgType=5) with tag 58-Text = **Logout initiated due to disconnect on primary connection** will be sent from CME Globex to the customer.
- If the primary gateway fails, failover occurs by promoting the backup gateway to assume the primary role.
 - During the designated primary instance failure condition, CME Group may keep the designated backup instance as primary gateway.
 - Customers who experience log on issues with their primary iLink sessions must rely on the reason text (tag 58) to detect if they are connecting to primary gateway, which could be different than the primary host IP address listed on the [CME SFTP -MSGW config.xml file](#).
 - To re-establish connection successfully, customer must log on to their backup iLink sessions.

i Client systems must examine tag 58-Text of the logout message to successfully reestablish connection to the newly promoted primary.

iLink Fault Tolerance Scenarios

Scenarios	Given Primary Gateway State	Given Backup Gateway State	When	Then Primary Connection State	Then Backup Connection State	Result
Initial Logon to Primary Session	Not connected	Not connected	Logon to primary followed by Logon to backup	Logged on	Logged on	Client system uses the fault tolerance feature. Client system must send a Logon message to primary gateway first with the Fault Tolerance Indicator (FTI) in tag 49-SenderCompID set to 'U' for undefined.

Initial Logon to Primary Session	Not connected	Not connected	Logon to backup		Logged out	New behavior Client system sends a Logon to backup gateway first. CME Globex sends a Logout (tag 35-MsgType=5) message.
Primary Session Disconnect	Logged on	Logged on	Disconnect from primary	Disconnected	Logged out	Client system is disconnected from primary gateway, the backup connection is logged out as well. All resting orders are cancelled for all Cancel on Disconnect registered iLink sessions. All resting quotes are cancelled on behalf of the Mass Quote customers.
Backup Session Disconnect	Logged on	Logged on	Disconnect from backup	Logged on	Disconnected	A TCP disconnection from backup gateway does not affect primary connection. Resting orders and quotes are not cancelled.
Logout from Primary Session	Logged on	Logged on	Logout from primary	Logged out	Logged out	Client system is logged out from primary gateway, the backup connection is logged out as well. If the logout is exchange initiated, all resting orders are cancelled for all Cancel on Disconnect registered gateway sessions. All resting quotes are cancelled on behalf of the Mass Quote customers.
Logout from Backup Session	Logged on	Logged on	Logout from backup	Logged on	Logged out	Logout from backup gateway does not affect primary connection. All resting orders and quotes are not cancelled.
Logon with Fault Tolerance (FTI value of U, P, or B)	Logged on	Logged on	Primary MSGW Failure	Disconnected	Backup connection becomes Primary	Newly promoted primary will send a heartbeat. Resting orders are not cancelled. All resting quotes are cancelled on behalf of Mass Quote customers.
Logon with Fault Tolerance (FTI value of U, P, or B)	Logged on	Not connected	Primary MSGW Failure	Disconnected	Session must connect and login to the new promoted primary.	All resting orders are cancelled for all Cancel on Disconnect registered iLink sessions. All resting quotes are cancelled on behalf of Mass Quote customers.
Logon without Fault Tolerance (FTI value of N)	Logged on	Not allowed to connect	Primary MSGW Failure	Disconnected	Session must connect and login to the new promoted primary.	All resting orders are cancelled for all Cancel on Disconnect registered iLink sessions. All resting quotes are cancelled on behalf of Mass Quote customers.
Logon with Fault Tolerance (FTI value of U, P, or B)	Logged on	Logged on	Backup MSGW Failure	Logged on	Session must connect and login to the restored backup.	The failed backup MSGW is restored. All resting orders and quotes are not cancelled.
Logon without Fault Tolerance (FTI value of N)	Logged on	Not allowed to connect	Backup MSGW Failure	Logged on	Not connected	The status of the primary connection remains intact. All resting orders and quotes are not cancelled.



- A TCP disconnection or a FIX logout can be either client-initiated or Exchange-initiated.
- A TCP disconnection or a FIX logout does not lead to failover; only gateway failure leads to failover.
- When client systems are logged out or disconnected from the primary connection, they will be logged out from the backup connection as well, if applicable. (Does not apply to the Primary gateway failure scenario.)

Convenience Gateway

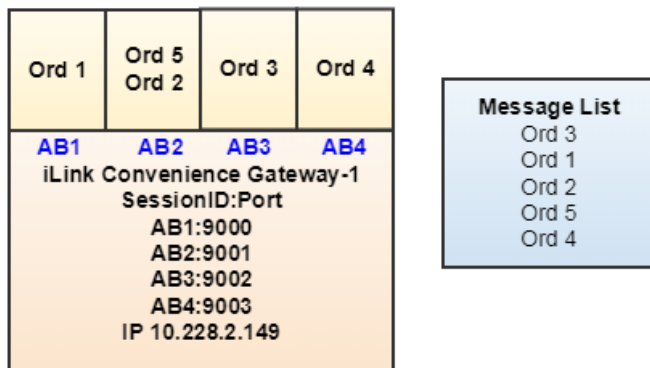
Customers can continue access and trade on the "standard" iLink gateway, referred to as the iLink Convenience Gateway (CGW).

The CGW allows the client system to connect to any CGW session, which manages routing for all order messages.

CGW FIFO Message Processing Model

The CGW FIFO message handling is enforced once the message reaches the MSGW. The following diagram shows the CGW model in which the example Session, AB1, exists on a single gateway.

- The IP address and port define a socket.
- An arriving message triggers an event for the gateway to read from its sockets.
- If numerous sockets have messages, the entire list is provided to the application.
- This list will not be ordered based upon message arrival



CGW Sequence Number Management

CGW message sequencing is managed using the SenderCompID value as in standard iLink [Session Layer - Message Sequence Numbers](#).

Market Segment Gateway

One iLink MSGW provides connectivity to a single market segment. Currently there are 17 market segments or groups of products on a given trading engine instance.

Client system messaging must identify the market segment for order entry routing. Market Segment is defined for each instrument in tag 1300-MarketSegmentID in the [MDP Security Definition message](#).

SFTP - MSGW Configuration


An SFTP site (sftpng.cmegroup.com), accessed via a CME Globex network direct connection, is used to disseminate the MSGW configuration information. The MSGW Configuration allows clients system to receive the list of all market segments and MSGW IPs. This SFTP site contains the configuration files for all environments. The SFTP site is a secure site that requires a user name and password for access.

Table Key:

- Environment - specific environment (i.e., Certification, New Release, Production).
- Note:** The Autocert+ tool will use the NR MSGW config file.
- Service - the Configuration service.
 - SFTP Site - address of SFTP site.
 - User Name - identifies the user name.
 - Password - identifies the password.
 - Directory Location - identifies the directory.
 - Client System Update Schedule - Client systems should download updates according to the schedule specified.

Environments	Service	SFTP Site	User Name	Password	Directory	Client System Update Schedule
--------------	---------	-----------	-----------	----------	-----------	-------------------------------

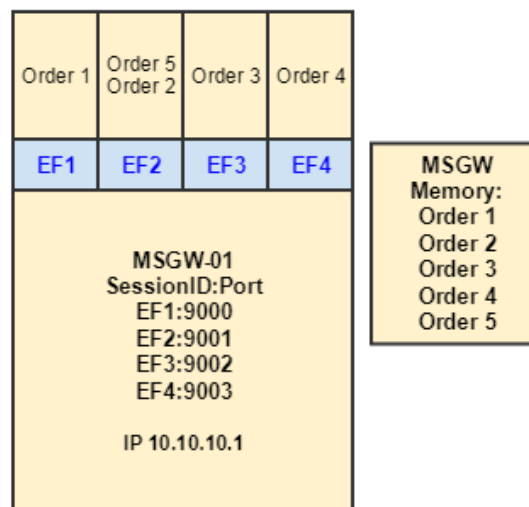
Production	Configuration	sftpng.cmegroup.com	cmeconfig	G3t (Onnect3d)	/MSGW/Production /Configuration	daily
Certification					/MSGW/Cert/Configuration	daily
New Release Certification					/MSGW/NRCert/Configuration	daily


 In addition to the generic User Name/Password, client systems can connect using the same credential currently used for CME SFTP site. Additional information pertaining to the CME Secure SFTP site is available in [CME Clearing Advisory Notice 15-105](#).

MSGW FIFO Message Processing

The MSGW ensures messages are processed in the order in which they arrive at the gateway.

Arriving message is immediately placed in MSGW memory; messages are processed in the order of arrival at the gateway.




 Packet size maximum behavior based on MTU (Maximum Transmission Unit) - 1420 bytes, as defined on our network

- If a customer sends a packet with an incomplete message, the message will be processed once the subsequent packet completes the message. The message is entered into the FIFO queue on the MSGW only when the message is complete.
- If a customer sends a packet with one or more complete messages and a partial message at the end, all complete messages from that packet will be processed immediately and the partial message is held to be processed on the subsequent packet(s).
- If a customer sends a packet with an incomplete message and the first complete message of the next packet on that session is not the continuation of the message, the MSGW will merge both messages into a single message and then discard the message due to the invalid tag 10-Checksum value.

MSGW Sequence Number Management

A single session is able to login to all 17 MSGW instances. To manage message sequence numbers, the customer must leverage the combination of Session ID (the first 3 characters of tag 49-SenderCompID) and tag 1300-MarketSegmentID.

CME Globex is currently comprised of 17 market segments. The number of market segments and the instruments listed on each can change in the future. **Customers are strongly encouraged to rely programmatically on tag 1300-MarketSegmentID in the Market Data Security Definition message to route orders to the correct MSGW.**

 Customers must determine how the new sequencing stream may impact order management systems on the customer side.

	CGW Session ID	MSGW Session ID
Number of TCP connection(s)	1 TCP connection to access all market segments	1 to 17 TCP connections to access all market segments
Number of sequence stream(s)	1 sequence stream	Customers must manage 1 to 17 different sequence streams based on the combination of Session ID and Market Segment ID

Tag 57-TargetSubID

Customers may connect directly to all Market Segment Gateways using the same Session ID. Customers can use tag 1300-MarketSegmentID from the market data Security Definition (tag 35-MsgType=d) message to specify the market segment for the products they want to trade on CME Globex.

For iLink messages sent from the client system to CME Globex, MSGW customers can leverage tag 57-TargetSubID to specify the market segment of the instrument. Tag 57-TargetSubID contains the tag 1300-Market SegmentID value. CME Globex will respond and echo back tag 57-TargetSubID values submitted in the customer iLink message on tag 50-SenderSubID.

Example

Time	From Client System to CME Globex	From CME Globex to Client System	Tag 35-MsgType	Tag 49-SenderCompID	Tag 50-SenderSubID	Tag 57-TargetSubID
1	Logon		A	ZUB007N	ATS	74
2		Logon Ack	A	CME	74	ATS
3		Test Request	1	CME	74	ATS
4	New Order		D	ZUB007N	ATS	74
5		Execution Report	8	CME	74	ATS
6	Logon		A	ZUB007N	ATS	68
7		Logon Ack	A	CME	68	ATS
8		Test Request	1	CME	68	ATS
9	New Order		D	ZUB007N	ATS	68
10		Execution Report	8	CME	68	ATS

- MSGW will reject any attempts to logon with an incorrect Market Segment ID with a Logout message (tag 35=5) including tag 58 = **Invalid TargetSubID(tag 57)**.
- MSGW will reject any application message with an instrument not assigned to that particular Market Segment ID with a Session Level Reject (tag 35=3) including tag 58 = **Invalid SecurityDesc(tag 107)**.



CGW sessions will continue to use tag 57-TargetSubID=G for all messages sent from the client system to CME Globex.

MSGW Failover Scenarios

During a MSGW failover from the primary to the backup instance, customer iLink messages will be rejected.



The GCC does not announce MSGW failovers, since the information is already available real-time to all market participants via MDP 3.0 and iLink session behavior. In the event of a MSGW failover, MSGW session customers must follow the recovery process outlined here. Any MSGW failover that does not follow standard behaviors, or for which market data is not available, will be announced.

During the primary MSGW failover to backup:

- CME Globex initiates failover to backup MSGW.
 - CGWs will reject all incoming business messages for instruments on the impacted Market Segment with a [Session Level Reject \(tag 35=3\)](#) with tag 58=**Technical Error – Market Segment: <MarketSegmentID> is not available at this time**.
- If the MSGW failing over is for a futures engine, all of the options on the associated options engine will automatically transition to a trading halt ([MDP Security Status\(tag 35-MsgType=f\) message](#) with tag 326-SecurityTradingStatus=2 and tag 327-HaltReason=0). This trading halt applies to **all** of the options on the market segment. All resting mass quotes are cancelled during the transition. During the halt, customers can submit, modify and cancel orders. Mass quotes are rejected. This halt is maintained for a few seconds after the futures MSGW has resumed trading.
- **Technical Pause** is introduced, which allows customers to cancel existing orders and to ensure that all messages sent during the MSGW failover are acknowledged by CME Globex.
 - Technical Pause is indicated by [MDP Security Status \(tag 35-MsgType=f\) message](#) with:
 - tag 326-SecurityTradingStatus=2 (Trading Halt) and tag 327-HaltReason=6 (Recovery in process)

- During the Technical Pause **only order cancellation** is allowed.
- There are 2 scenarios where technical pause will be initiated:
 - The current market state is in a less restrictive state than Pause (i.e. Open, Pre Open)
 - The current market state is already in Pause

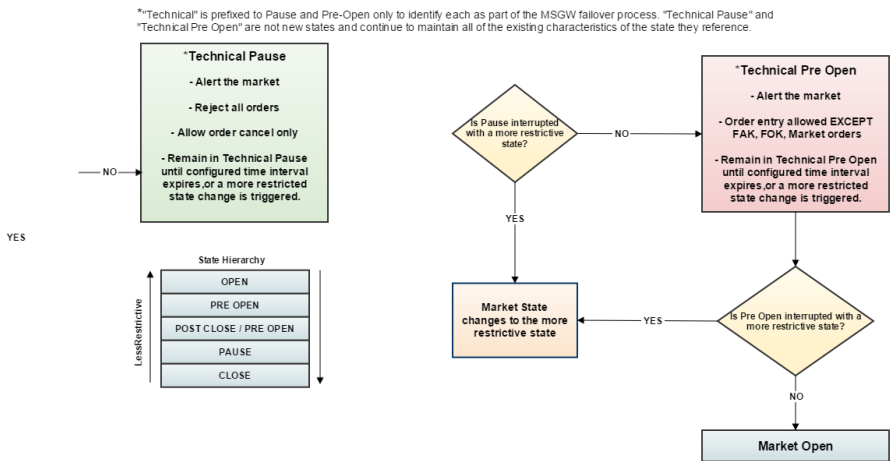
i Technical Pause is effective for a configurable duration—currently 10 seconds—before the market transitions to Technical Pre Open. If, at any point during the Technical Pause, CME Globex receives a more restrictive market state change, it will discontinue Technical Pause and enter the more restrictive state.

- After the Technical Pause:
 - MSGW will replay customer iLink messages that were not processed during the failover
 - CME Globex will reject those replayed messages, except Cancel messages.
- **Technical Pre Open** - Technical Pre Open will follow the technical pause to allow adequate time for order entry and price discovery before reopening the markets.
 - Technical Pre Open MDP Security Status(tag 35-MessageType=f) message is generated with:
 - MDP 3.0: tag 326-SecurityTradingStatus=21 (Pre Open) and tag 327-HaltReason=6 (Recovery in process).
 - tag 326-SecurityTradingStatus=21 (Pre Open)
 - During the Technical Pre Open state, Order Entry is allowed, except for FAK, FOK, and Market order types.
 - Similar to the Technical Pause, if the current market state is more restrictive than Pre Open, Pre Open will not be initiated.

i Technical Pre Open is effective for a configurable duration—currently 10 seconds—before market transitions to opening. If, at any point during the Technical Pre Open, CME Globex receives a more restrictive market state change, it will stop Technical Pre Open and enter the more restrictive state

- After the market is re-opened, order entry will be allowed.

The following diagram illustrates the process above.



EXCEPTION
 Session Level Reject (tag 35=3) with tag 58=Technical Error – Market Segment: <MarketSegmentID> is not available at this time is sent for CGW sessions only.

Backup MSGW 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.

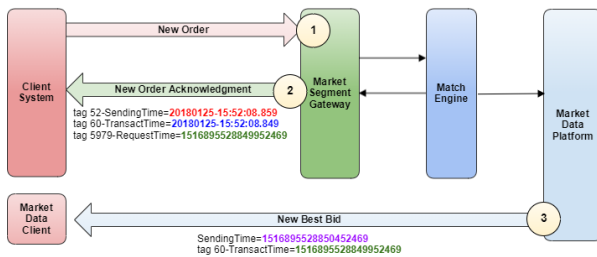
CME Globex Timestamps

Watch this brief video for an overview of CME Globex timestamps.



The diagram shows all timestamps for iLink and MDP 3.0 messages:

- All MDP 3.0 timestamps are in nanoseconds with microsecond precision.
- The iLink timestamp (tag 5979-RequestTime) is in nanoseconds with microsecond accuracy.
- Timestamp(1) on iLink and MDP 3.0 represents the time when the order reaches the Market Segment Gateway.
- Timestamp(1) is included in the MDP 3.0 Trade Summary and Order Book Update messages.
- A customer can compare timestamps with an aggressor's timestamp to determine how much later the customer's order was submitted.



- 1 Tag 60-TransactTime - hardware timestamp generated when incoming message hits Market Segment Gateway (UTC format). Sent to client in outbound message.
Tag 5979-RequestTime - hardware timestamp generated when incoming message hits Market Segment Gateway (in nanosecond format with microsecond accuracy). Sent to client in outbound message and in market data in tag 5979-RequestTime.
- 2 Tag 52-SendingTime - software timestamp generated when outbound message to client leaves the Market Segment Gateway.
- 3 SendingTime - MDP3.0 Binary Header software timestamp generated when the message leaves the market data gateway.
Tag 60-TransactTime contains the equivalent iLink tag 5979 value from step 1 in UTC format with nanosecond precision.

Nanosecond Timestamp (tag 5979-RequestTime)

To provide more transparency between iLink order entry and MDP 3.0 market data, iLink supports a FIX tag 5979-RequestTime on all Convenience Gateway (CGW) and Market Segment Gateway (MSGW) sessions.

This timestamp represents the time the customer iLink message was received by the Market Segment Gateway and had FIFO order handling applied. The tag provides the date and time, in Universal Time Coordinated (UTC), represented in nanoseconds past the UNIX epoch time (00:00:00 UTC on 1 January 1970). UTC Timestamps are sent in number of nanoseconds since Unix epoch synced to a master clock to microsecond accuracy.

CME Globex uses SolarFlare's WODA (wire order delivery API) for this new timestamp.

Tag	Name	Req	Type	Max Length	Description
-----	------	-----	------	------------	-------------

5979	Request Time	N	Int	20	Information carried on a response to convey the time (UTC) the outbound message was processed by the MSGW application. UTC Timestamps are sent in number of nanoseconds since Unix epoch synced to a master clock to microsecond accuracy.
------	--------------	---	-----	----	--

Tag 5979-RequestTime will be added to the following iLink application messages sent from CME Globex to the client system:

- Execution Report: Order Creation, Cancel, or Modify (tag 35-MsgType=8, tag 39-OrdStatus=0, 4 or 5)
- Execution Report: Fill Notice (tag 35-MsgType=8, tag 39-OrdStatus=1 or 2)
- Execution Report: Trade Cancel (tag 35-MsgType=8, tag 39-OrdStatus=H)
- Execution Report: Order Elimination (tag 35-MsgType=8, tag 39-OrdStatus=4 or C)
- Execution Report Order Status Request Acknowledgment (tag 35-MsgType=8, tag 150-ExecType=I)
- Execution Report: Reject (tag 35-MsgType=8, tag 39-OrdStatus=8)
- Order Cancel/Replace Reject (tag 35-MsgType=9, tag 434-CancelRejResponseTo=2)
- Order Cancel Reject (tag 35-MsgType=9, tag 434-CancelRejResponseTo=1)
- Order Mass Action Report (tag 35-MsgType = BZ)
- Quote Acknowledgement (tag 35-MsgType=b)
- Quote Reject (tag 35-MsgType=b, tag 297-QuoteAckStatus =5)
- Quote Cancel Acknowledgement (tag 35-MsgType=b, tag 297-QuoteAckStatus=1, 3, 4 or 5 without tag 9775-UnsolicitedCancelType)
 - Only for Solicited Quote Cancel

All Administrative messages and the following Application messages sent from CME Globex to the client system will not include this new tag:

- Quote Cancel Acknowledgment (tag 35-MsgType=b, Tag 9775--UnsolicitedCancelType=A, B, C, D, E, F or H)
 - Only for Unsolicited Quote Cancel



To ensure the 5979 timestamp never reflects an out-of-sequence event, in certain cases CME Globex will increment timestamp value. These forced timestamps still accurately reflect the sequence of events on the MSGW.

iLink Unsolicited Outbound Message Scenario

For unsolicited outbound messages where there are no iLink inbound messages, tag 5979 will represent the time at which the transaction is processed on the trading engine in nanoseconds.

Scenario	Tag 5979	Tag 60
Fill Execution Reports (35=8, 39=1,2)	The time at which the aggressing order reached the MSGW in nanoseconds. If no aggressor (e.g., Pre-open), will contain the engine event time in nanoseconds.	The time at which the aggressing order reached the MSGW in milliseconds If no aggressor (e.g., pre-open), contains the engine event time in milliseconds.
Stop Order Trigger Ack (35=8, 39=0, 40=2)	The time at which the original aggressing order, which initiated the fill event that triggered the stop order, reached the MSGW in nanoseconds. If no aggressor (e.g., Pre-open) is present, will contain the engine event time in nanoseconds.	The time at which the original aggressing order, which initiated the fill event that triggered the stop order, reached the MSGW in milliseconds. If no aggressor (e.g., Pre-open) is present, will contain the engine event time in milliseconds.
FAK Elimination Acks (35=8, 39=C)	The time at which the corresponding order reached the MSGW in nanosecond	The time at which the corresponding order reached the MSGW in millisecond
Trade Bust Execution Report (35=8, 39=H)	The time of trade bust event in nanoseconds.	The time of the trade bust event in milliseconds.
Unsolicited Mass Quote Cancel Ack (35=b, 9775=A, B,C,D,E,F,G,H)	Not present.	Not present.
Unsolicited Cancel Ack (35=8, 39=4, 378=8, 100,103,104,105,106,107)	The engine event time in nanoseconds.	The engine event time in milliseconds.
End of Day Elimination Ack (35=8, 39=C)	The time of the order elimination event in nanoseconds.	The time of the order elimination event in milliseconds.

Tag 60-TransactTime contains the time when a customer iLink message reaches the MSGW, where applicable. Otherwise, tag 60 represents the trading engine event time.

Customers may use iLink tag 60 or tag 5979 interchangeably, depending on the level of timestamp granularity needed, and can leverage them in combination with the MDP 3.0 tag 60 value.

- the iLink tag 60 value will be derived from iLink tag 5979-RequestTime, but represented in the existing UTC format with millisecond precision.
- MDP 3.0 tag 60 timestamp is sent in number of nanoseconds since the Unix epoch synced to a master clock to microsecond accuracy.

Testing and Certification

An iLink MSGW certification suite is available in Autocert+. Both iLink and the iLink MSGW certifications are required for customer systems to connect directly to the MSGW. To certify, customers must use the same System Profile for both iLink and iLink MSGW certifications.

New Release Weekly Failover Testing Windows

Market Segment ID (tag 1300)	Market Segment Description	Date	Time
54	CME Equity options	Tuesday and Thursday	11:00 pm CT - MSGW failover to backup component 11:30 pm CT - MSGW failover to primary component
70	CME Commodity Futures	Tuesday and Thursday	9:00 am CT - MSGW failover to backup component 9:30 am CT - MSGW failover to primary component
84	CBOT Treasury Futures	Tuesday and Thursday	11:00 pm CT - MSGW failover to backup component 11:30 pm CT - MSGW failover to primary component