

T7 Release 9.0

Incident Handling Guide

Version	4.7
Status	Final
Date:	31-08-2020

© 2020 Copyright by Deutsche Börse AG ("DBAG"). All rights reserved.

All intellectual property, proprietary and other rights and interests in this publication and the subject matter of this publication are owned by DBAG or other entities of Deutsche Börse Group. This includes, but is not limited to, registered designs and copyrights as well as trademark and service mark rights.

Specifically, the following trademarks and service marks are owned by entities of Deutsche Börse Group:

Buxl®, DAX®, DivDAX®, eb.rexx®, Eurex®, Eurex Bonds®, Eurex Repo®, Eurex Strategy WizardSM, Euro GC Pooling®, F7®, FDAX®, FWB®, GC Pooling®, GCPI®, M7®, MDAX®, N7®, ODAX®, SDAX®, T7®, TecDAX®, USD GC Pooling®, VDAX®, VDAX-NEW® and Xetra® are registered trademarks of DBAG.

The following trademarks and service marks are used by Deutsche Börse Group under license and are property of their respective owners:

All MSCI indexes are service marks and the exclusive property of MSCI Barra. ATX®, ATX® five, CECE® and RDX® are registered trademarks of Vienna Stock Exchange AG.

IPD® UK Annual All Property Index is a registered trademark of Investment Property Databank Ltd. IPD and has been licensed for the use by Eurex for derivatives.

SLI®, SMI® and SMIM® are registered trademarks of SIX Swiss Exchange AG.

The STOXX® indexes, the data included therein and the trademarks used in the index names are the intellectual property of STOXX Limited and/or its licensors Eurex derivatives based on the STOXX® indexes are in no way sponsored, endorsed, sold or promoted by STOXX and its licensors and neither STOXX nor its licensors shall have any liability with respect thereto. Bloomberg Commodity IndexSM and any related sub-indexes are service marks of Bloomberg L.P.

PCS® and Property Claim Services® are registered trademarks of ISO Services, Inc.

Korea Exchange, KRX, KOSPI and KOSPI 200 are registered trademarks of Korea Exchange Inc. Taiwan Futures Exchange and TAIFEX are registered trademarks of Taiwan Futures Exchange Corporation. Taiwan Stock Exchange, TWSE and TAIEX are the registered trademarks of Taiwan Stock Exchange Corporation.

BSE and SENSEX are trademarks/service marks of Bombay Stock Exchange (BSE) and all rights accruing from the same, statutory or otherwise, wholly vest with BSE. Any violation of the above would constitute an offence under the laws of India and international treaties governing the same.

Methods and devices described in this publication may be subject to patents or patent applications by entities of Deutsche Börse Group.

Information contained in this publication may be erroneous and/or untimely. Neither DBAG nor any entity of Deutsche Börse Group makes any express or implied representations or warranties regarding the information contained herein. This includes any implied warranty of the information's merchantability or fitness for any particular purpose and any warranty with respect to the accuracy, correctness, quality, completeness or timeliness of the information.

Neither DBAG nor any entity of Deutsche Börse Group shall be responsible or liable for any errors or omissions contained in this publication, except for DBAG's or the respective Deutsche Börse Group entity's wilful misconduct or gross negligence.

Neither DBAG nor any entity of Deutsche Börse Group shall be responsible or liable for any third party's use of any information contained in this publication under any circumstances.

All descriptions, examples and calculations contained in this publication are for illustrative purposes only, and may be changed without further notice.

Content

1. Introduction	6
1.1 Content of this document	6
1.2 Usage Notes	6
1.3 Further reading	6
2. Concepts	7
3. Scenarios	9
3.1 Low Frequency Gateway Failure	9
3.1.1 ETI Interface	9
3.1.2 Market Data – Enhanced Orderbook Interface (EOBI)	9
3.1.3 Market Data – Enhanced Market Data Interface (EMDI)	9
3.1.4 GUI	9
3.1.5 Best Practice	9
3.2 Partition Specific Gateway Failure	9
3.2.1 ETI Interface	10
3.2.2 Market Data – Enhanced Orderbook Interface (EOBI)	10
3.2.3 Market Data – Enhanced Market Data Interface (EMDI)	10
3.2.4 GUI	10
3.2.5 Best Practice	10
3.3 Low Frequency Gateway losing connection to Partition Specific Gateway	10
3.3.1 ETI Interface	10
3.3.2 Market Data – Enhanced Orderbook Interface (EOBI)	11
3.3.3 Market Data – Enhanced Market Data Interface (EMDI)	11
3.3.4 GUI	11
3.3.5 Best Practice	11
3.4 Stalled Gateway	11
3.4.1 ETI Interface	11
3.4.2 Market Data – Enhanced Orderbook Interface (EOBI)	11
3.4.3 Market Data – Enhanced Market Data Interface (EMDI)	11
3.4.4 GUI	11
3.4.5 Best Practice	11
3.5 Matching Engine failover	12
3.5.1 ETI Interface	12

3.5.2	Market Data– Enhanced Orderbook Interface (EOBI)	12
3.5.3	Market Data – Enhanced Market Data Interface (EMDI)	13
3.5.4	GUI	13
3.6	Matching Engine Processing Delay	13
3.6.1	Order Entry (ETI)	13
3.6.2	Market Data – Enhanced Orderbook Interface (EOBI)	13
3.6.3	Market Data – Enhanced Market Data Interface (EMDI)	13
3.6.4	GUI	14
3.6.5	Best Practice	14
3.7	Market Data failover – MDI, EMDI and EOBI	14
3.7.1	Order Entry (ETI)	14
3.7.2	Market Data – Enhanced Orderbook Interface (EOBI)	14
3.7.3	Market Data – Market Data Interface (MDI)	14
3.7.4	Market Data – Enhanced Market Data Interface (EMDI)	14
3.7.5	GUI	14
3.8	Persistency layer failover	15
3.9	Central component failure	15
3.9.1	Matching Engine	15
3.9.2	Market Data – Enhanced Orderbook Interface (EOBI)	16
3.9.3	Market Data – Enhanced Market Data Interface (MDI, EMDI)	16
3.9.4	GUI	16
3.9.5	Best Practice	16
3.10	T7 Entry Service failover and failure	18
3.10.1	Order Entry (ETI)	18
3.10.2	Market Data – Enhanced Orderbook Interface (EOBI)	18
3.10.3	Market Data – Enhanced Market Data Interface (EMDI)	18
3.10.4	GUI	18
3.10.5	Best Practice	18
3.11	GUI unavailability	18
3.11.1	Order Entry (ETI)	18
3.11.2	Market Data – Enhanced Orderbook Interface (EOBI)	19
3.11.3	Market Data – Enhanced Market Data Interface (EMDI)	19
3.11.4	GUI	19
3.11.5	Best Practice	19

4. Appendix	20
4.1 Handling of Matching Engine Service Availability notifications	20
4.1.1 Best practice	20
5. Change log across the Releases	21

1. Introduction

1.1 Content of this document

This document provides a detailed description of the reaction of the T7 trading system to technical incidents. It describes the reaction of T7 in scenarios either observed in the past (e.g. Matching Engine Failover) or to be expected in case of technical outages (e.g. Market Data failover).

The *Market Status Indicator* on the Eurex and Xetra websites (<https://www.eurexchange.com/> / <https://www.xetra.com>) can serve as an early indicator for any major technical issues in the T7 system.

1.2 Usage Notes

Even though care has been taken to describe and illustrate known scenarios, the list may not be complete as yet unobserved failure modes might occur in the future

1.3 Further reading

The following T7 documents provide additional information to complement the description given in this document:

- *T7 Functional and Interface Overview*
- *T7 Eurex Trader and Admin GUI User Manual*
- *T7 Cash Markets Trader, Admin and Clearer GUI – Manual*
- *T7 Enhanced Trading Interface – Manual*
- *T7 Market and Reference Data Interfaces – Manual*
- *T7 Extended Market Data Service – Manual*
- *T7 Enhanced Order Book Interface – Manual*
- *T7 Network Access Manual*

Some of the scenarios referred to in this document are also described in the Participant Simulation Guide and the days when they take place is indicated in the Simulation Calendar.

These and other documents providing information for the T7 trading system can be found on the web sites of the respective exchanges, i.e.

<https://www.eurexchange.com/exchange-en/technology/t7/release9>

and

<https://www.xetra.com/xetra-en/technology/t7/system-documentation/release9>

2. Concepts

The T7 trading system is designed for highest availability and robustness. A guiding principle has been to architect the system in such a way that it can cope with failures of individual components with the minimal effect on the functioning of the market. This is achieved on the one hand by separating the market services into completely independent groups of products (partitions) and on the other hand via technical failover mechanisms.

Central components (e.g. Matching Engine, Persistency layer, T7 Entry Service) have a standby partner that takes over immediately when the primary component fails. In most cases this **failover** is realized transparently, i.e. without visible impact on any of the T7 interfaces. In case of a non-transparent failover the visible effect on the affected interfaces is described in the following sections.

The T7 system is not prepared to automatically handle the simultaneous failure of primary and standby partner – a double failure of a central component. The component shows then a complete **failure**. Such incidents are not resolved transparently and require manual interception from T7 operations. The scenarios have a longer impact on interfaces and service availabilities than transparent or non-transparent failover scenarios. In such situations as a general advice, members should follow the exchanges Newsboard messages via ETI News Broadcast, T7 Trading GUI Newsboard or the Newsboard Webpage, i.e.

<https://www.eurexchange.com/exchange-en/trading/production-newsboard>

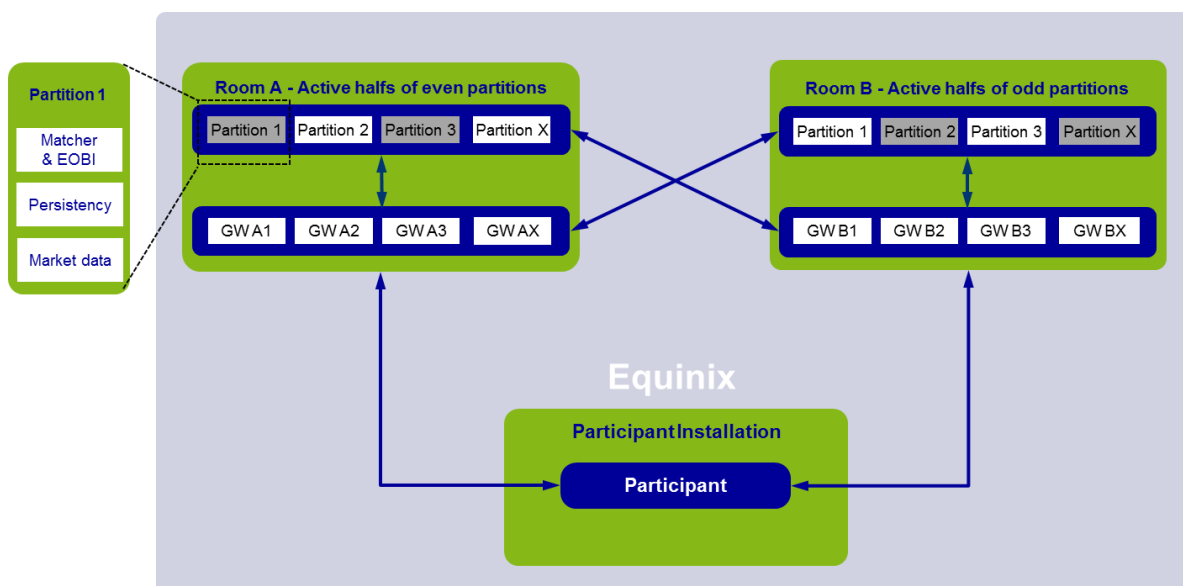
<https://www.xetra.com/xetra-en/newsroom/xetra-newsboard>

The member portal provides an additional SMS and Email subscription service to Emergency Notifications which are sent in cases of more disruptive market events and major failures like a complete T7 system failure.

<https://member.deutsche-boerse.com> → My Profile → Emergency Notification

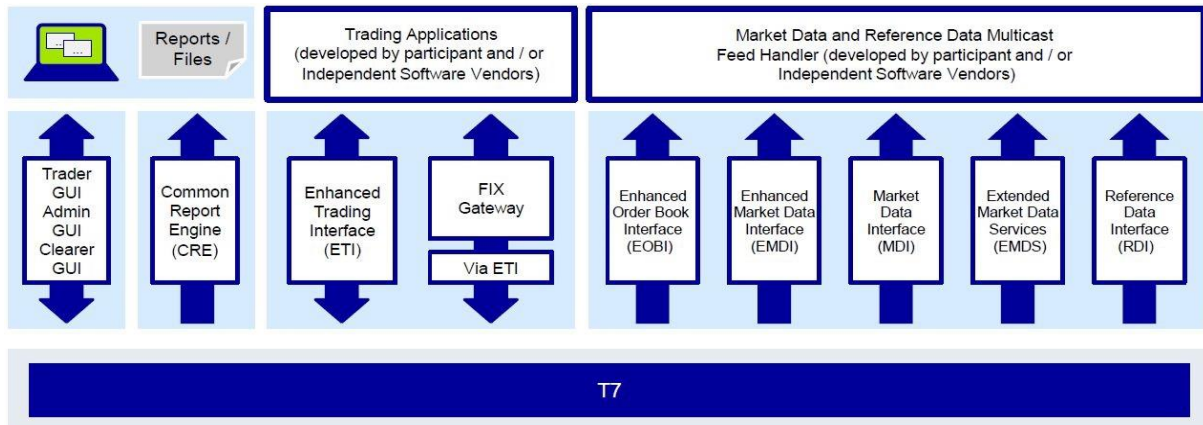
As a last resort call Trading Operations (Derivatives) at +49-69-211-11210/Cash Markets Operations (+49-69-211-11400).

The concept of two redundant backend rooms also guarantees service availability in case a complete room in the data center has a major technical outage. The primary and secondary partners of central processing components are always deployed in different rooms to ensure full exchange functionality in case of a room outage. The schematic overview of this principle is shown in the following picture:



The failover of a T7 Gateway is not transparent, i.e. a failure of a Gateway always requires member interaction, usually reconnecting an affected trading session to a different Gateway. This ensures exchange connectivity even in the case a member loses one of its network connections.

The remainder of this document describes failover and failure scenarios and how they materialize on the T7 Interfaces depicted in the overview below. If not stated otherwise the scenarios describe the behavior of market services for one partition.



The FIX Gateway and the Extended Market Data Service interfaces are not covered within this document.

The RDI interface has the same technical interface and behavior as the MDI interface for all described scenarios and is for the remainder of this document not explicitly mentioned.

The Common Report Engine interface is not affected by any of the described scenario and thus not included in any of the descriptions.

3. Scenarios

The following section describes the individual scenarios and the reaction of the T7 trading system mainly from an interface point of view.

3.1 Low Frequency Gateway Failure

A low frequency (LF) Gateway service might become unavailable for reasons such as hardware failure, network outages or software problems.

In the event of an LF Gateway failure, active sessions will be disconnected. The failure of the Gateway component will be detected, and all non-persistent orders and quotes of sessions logged in via the failed Gateway will be deleted¹ by the T7 system.

If this LF Gateway is connected to partition specific (PS) Gateway (see section 3.2) and the PS Gateway detects the LF Gateway failure earlier than the T7 core system the MassActionReason code in MassActionReports will contain "100 = Internal connection loss" instead of "6 = Session loss or logout".

3.1.1 ETI Interface

A failure of a Gateway leads to the termination of the TCP connection and a reconnection to the same Gateway is impossible.

3.1.2 Market Data – Enhanced Orderbook Interface (EOBI)

There is no special message for a failure of a Gateway. However as all non-persistent orders and quotes from sessions connected via the failed Gateway will be deleted, a lot of *Order Delete* updates are expected.

3.1.3 Market Data – Enhanced Market Data Interface (EMDI)

There is no special message for a failure of a Gateway. However as all non-persistent orders and quotes from sessions connected via the failed Gateway will be deleted, multiple price levels may see a reduction of size.

3.1.4 GUI

A failure of a Gateway is not indicated via the GUI. Secondary effects, i.e. order deletions of non-persistent standard orders will be visible in the order views.

3.1.5 Best Practice

If an active ETI session is disconnected, for any reason, the participant application should reconnect to another low frequency gateway (IP addresses are listed in the network access guide).

3.2 Partition Specific Gateway Failure

A PS Gateway service might become unavailable for reasons such as hardware failure, network outages or software problems. Active sessions on that gateway will be disconnected and all non-persistent orders and quotes of sessions logged in via the failed PS Gateway will be deleted.

¹ In case a Gateway failure is not (immediately) detected see Section 3.3

Stalled Gateway

The designated standby PS Gateway will be activated and will start accepting session logins. After a timeout of 30 seconds (configured value at time of this writing) order maintenance will become available on the now active PS Gateway.

3.2.1 ETI Interface

A failure of a Gateway leads to the termination of the TCP connection. The T7 System will recognize the failure of the PS Gateway and will activate the standby PS Gateways. Sessions may subsequently login to the activated PS Gateway. T7 broadcasts a message of the form "<mic>: Partition x: PS Gateway failover triggered, please login to standby PS Gateway".

3.2.2 Market Data – Enhanced Orderbook Interface (EOBI)

There is no special message for a failure of a PS Gateway. However as all non-persistent orders and quotes from sessions connected via the failed Gateway will be deleted, a lot of *Order Delete* updates are expected.

3.2.3 Market Data – Enhanced Market Data Interface (EMDI)

There is no special message for a failure of a PS Gateway. However as all non-persistent orders and quotes from sessions connected via the failed Gateway will be deleted, multiple price levels may see a reduction of size.

3.2.4 GUI

A failure of a Gateway is not indicated via the GUI. Secondary effects, i.e. order deletions of non-persistent standard orders will be visible in the order views.

3.2.5 Best Practice

If an active ETI session is disconnected, for any reason, the participant application should first try to reconnect to the PS Gateway and in case of a failed connection should then try the secondary PS Gateway (IP addresses available in the network access guide). The Session Login Response (or reject) will inform about the status of the PS Gateway.

3.3 Low Frequency Gateway losing connection to Partition Specific Gateway

All LF Gateway requests are routed via the corresponding PS Gateway to the respective target partition. In case the connection between a LF Gateway and a PS Gateway fails the T7 system deletes all non-persistent orders and quotes for affected sessions and partition and the *MassActionReason* code in *MassActionReport* broadcasts is set to the "100 = *Internal connection loss*".

The connection may fail due to PS Gateway failover (see section 3.1) or due to a technical issue between a single LF Gateway and PS Gateway.

3.3.1 ETI Interface

A **ServiceAvailabilityBroadcast** message indicates the *MatchingEngineStatus* on partition <x> is unavailable. The data below is for partition 1 (FDAX)

<pre>ServiceAvailabilityBroadcast.PartitionID = 1 ServiceAvailabilityBroadcast.MatchingEngineStatus = 0</pre>

Incoming messages for this service are from now on rejected directly in the Gateway. If the loss is not caused by a PS Gateway Failure/Failover other Low Frequency Gateways may still be able to route messages to this service.

The T7 system deletes all non-persistent orders and quotes for sessions logged in on the Low Frequency Gateway on the affected partition.

When the connection has been re-established this will be signaled by another **ServiceAvailabilityBroadcast** message with availability set to 1 for the partition:

```
ServiceAvailabilityBroadcast.PartitionID = 1
ServiceAvailabilityBroadcast.MatchingEngineStatus = 1
```

3.3.2 Market Data – Enhanced Orderbook Interface (EOBI)

There is no special message for a failure of a LF Gateway to PS Gateway connection. However, all non-persistent orders and quotes from affected sessions will be deleted for the respective partition.

3.3.3 Market Data – Enhanced Market Data Interface (EMDI)

There is no special message for a failure of a LF Gateway to PS Gateway connection. However as all non-persistent orders and quotes from affected session will be deleted, multiple price levels may see a reduction of size.

3.3.4 GUI

There is no visible effect in the Trader or Admin GUI.

3.3.5 Best Practice

Use a failover/standby session to a different Low Frequency Gateway and in case of uncertainty about the partition state revert to the Best Practices of section 3.2.5.

3.4 Stalled Gateway

Parts of the Gateway might become inaccessible, e.g. stalled but the TCP connection is not closed or reset. In this case the auto deletion of non-persistent orders and quotes may be delayed. We therefore recommend taking actions to ensure the correct status of these orders as described in 3.4.5.

A two-way stalled TCP connection or internally locked up Gateway can be detected by missed ETI heartbeat messages from the Gateway.

A half open TCP connection (only the Gateway receive side does not work anymore) is arguably hard to detect since ETI heartbeats from the Gateway are still received. A Participant's application may detect this stalled connection by lack of responses within a reasonable time while other activity on other sessions or market data indicate a working T7 system for the respective product.

3.4.1 ETI Interface

In case of a two-way stalled TCP connection no ETI heartbeat messages are received from the Gateway but the TCP connection itself is not closed.

3.4.2 Market Data – Enhanced Orderbook Interface (EOBI)

There is no effect on EOBI data dissemination.

3.4.3 Market Data – Enhanced Market Data Interface (EMDI)

There is no effect on EMDI data dissemination.

3.4.4 GUI

There is no visible effect in the Trader or Admin GUI.

3.4.5 Best Practice

In case a stalled Gateway is suspected the TCP connection to it should be terminated and the according session should try to re-login. An application should not use the stalled Gateway but an alternative Gateway. If the session is still logged in because of the stalled Gateway the re-login attempt will trigger cancellation of all non-persistent orders (either lean or standard orders) and quotes for that session. If any orders or quotes were to be deleted, this event will be denoted with *MassActionReason (28721): 7 = "Duplicate Session Login"* in *Order/Quote Mass Cancellation Notification* via listener or session data.

If the above logon request is not answered the T7 Trading GUI can be used to delete pending orders and quotes in the *Risk Control* view to delete all orders (per user) or to delete all quotes (per session).

3.5 Matching Engine failover

In order to achieve the necessary throughput and latency requirements, the T7 Matching Engine handles all incoming orders and the state of the orders in non-persistent memory. The persistency layer receives the modification and persists all actions not marked as non-persistent, e.g. modification to persistent orders, all executions...

A matching engine failover is thus not transparent but always implies a market reset. In case the primary matching engine fails, the secondary component will take over by loading the persistent orders from the persistency layer and trading continues after a market reset - all non-persistent orders and quotes are deleted.

Note that some of the actions prior to a failure of the matching engine may have triggered a preliminary acknowledgement to the trading participant but may not have been persisted. It is thus crucial that the states of all orders are synchronized with the restatement of the orders after the Matching Engine recovered.

Please also refer to the section 'Matching Engine Failover and Failure' in the Simulation Guide.

3.5.1 ETI Interface

A **ServiceAvailabilityBroadcast** message indicates the MatchingEngineStatus on partition <x> is unavailable. The data below is for partition 1 (FDAX)

```
ServiceAvailabilityBroadcast.PartitionID = 1
ServiceAvailabilityBroadcast.MatchingEngineStatus = 0
```

Incoming messages for this service are from now on rejected directly in the Gateway. A market reset is issued via a TradingSessionStatusBroadcast message:

```
TradingSessionStatusBroadcast.RBCHeaderME.PartitionID = 1
TradingSessionStatusBroadcast.TradSesEvent = 102
```

The state of all products served by this partition will undergo the following product state transition:

- no transition if the product state was 'Start Of Day', 'Pre Trading', 'Trading', 'Post Trading'
- products will be set to 'Halt' if the product state was 'Closing'
- products will be set to 'Halt' under certain internal conditions (e.g. failover takes very long)

When the secondary matching engine took over this will be signaled by another **ServiceAvailabilityBroadcast** message with availability set to 1 for the partition:

```
ServiceAvailabilityBroadcast.PartitionID = 1
ServiceAvailabilityBroadcast.MatchingEngineStatus = 1
```

Trading then continues.

3.5.2 Market Data– Enhanced Orderbook Interface (EOBI)

The market reset is communicated via an *Order Mass Delete* message. This is then followed by *Order Add* messages of all persistent (non deleted) orders in the book.

Note that since EOBI is sent out from the matching engine EOBI failover will always happen when the matching engine fails over. A change of source IP of the affected multicast streams will happen as the secondary takes over. The *AppSeqNum* and *MsgSeqNum* in EOBI messages are both reset to 1 and the field *AppSeqResetIndicator* is set to indicate the failover.

3.5.3 Market Data – Enhanced Market Data Interface (EMDI)

The change in the order book is communicated via a *DepthIncremental* message. Usually there will be changes on multiple price levels because of the market reset.

3.5.4 GUI

Afore-mentioned *ServiceAvailability* changes are indicated in the GUI status bar and displayed in detail in the Newsboard window. Submitting orders for the affected partition during the failover time is not possible and requests will be rejected.

The change of availability is also reflected in a news broadcast of the form “<xmic> Partition <n> is down” and “<xmic> Partition <n> is technically up again, please check Newsboard for further information”.

The changed order book is displayed in corresponding market overviews and any product state change to “*Halt*” is indicated additionally in the Newsboard view.

3.6 Matching Engine Processing Delay

Large processing delays in the T7 matching engine (caused e.g. by technical problems or unexpected extreme loads) may lead to the built-up of internally queued orders. To avoid processing of old orders and to quickly recover from such a delay situation the affected products change to an internal “not tradeable” state. The incident is automatically detected and for the relevant products the following actions are taken:

- All non-persistent orders and quotes are deleted
- The product is temporarily set to an internal “not tradeable” state and session data and listener broadcasts are sent. The product is not set to “*Halt*”
- For a minimum time frame (currently configured to 10 seconds) and until the delayed processing is resolved all transactions will be rejected except for order deletions
- Once the situation is resolved this is communicated via session data and listener broadcasts

Please also refer to section ‘Matching Engine Processing Delay’ of the Simulation Guide.

3.6.1 Order Entry (ETI)

Deleted orders and quotes on entering the “not tradeable” state of a product are reported via session data. Additionally product-specific *DeleteAllOrderQuoteEventBroadcast* messages are received for the affected product(s) with *MassActionReason* set to *Product_temporarily_not_tradeable(11)*.

During the delay phase all transactions except deletions are rejected with *SessionRejectReason* set to *Service_Temporarily_Not_Available(102)* and *VarText* *TRANSACTION REJECTED DUE TO SLOW PARTITION*

When the matching engine will accept transactions again participants will be informed by a *TradingSessionStatus* message (*MsgType* (35) = “h”) specifying *TradSesEvent* (1368) = *Service Resume* (105)).

3.6.2 Market Data – Enhanced Orderbook Interface (EOBI)

There is no special message for a matching engine processing delay, but a lot of deleted orders and quotes are expected.

3.6.3 Market Data – Enhanced Market Data Interface (EMDI)

There is no indicator for a currently not tradeable product because of processing delays. Due to the deleted orders changes on multiple price levels are expected.

3.6.4 GUI

Order deletions are visible in the views for the affected products Order submissions for affected products will be rejected with *TRANSACTION REJECTED DUE TO SLOW PARTITION*.

3.6.5 Best Practice

If an ETI trading session receives a *Product_temporarily_not_tradeable* message (see 3.6.1) submission of new orders should be avoided until the *Service Resumed* notification is encountered since new order flow will only exacerbate the solution of the delay problem and the orders will be rejected anyway. Deletion of non persistent orders and quotes already took place so any clean-up can be limited to the deletion of persistent orders if necessary.

3.7 Market Data failover – MDI, EMDI and EOBI

In the event of a failure of a market data disseminator the standby component will take over as soon as the failure is detected. The failover will be completely transparent on all streams other than the failed stream.

Note that since EOBI is sent out from the matching engine an EOBI failover occurs only in case of a Matching Engine failover or failure.

Please also refer to section 'Matching Engine Failover and Failure, EOBI Failure' of the simulation guide and Section 'Failure of the market data feed/matching engine' in the T7 Market and Reference Data Interfaces Manual.

3.7.1 Order Entry (ETI)

No reaction to the failed market data components.

3.7.2 Market Data – Enhanced Orderbook Interface (EOBI)

The market reset is communicated via an *Order Mass Delete* message. This is then followed by *Order Add* messages of all persistent (non deleted) orders in the book.

Note that since EOBI is sent out from the matching engine EOBI failover will always happen when the matching engine fails over. A change of source IP of the affected multicast streams will happen as the secondary takes over. The *AppSeqNum* and *MsgSeqNum* in EOBI messages are both reset to 1 and the field *AppSeqResetIndicator* is set to indicate the failover.

3.7.3 Market Data – Market Data Interface (MDI)

A change of Source IP of the affected multicast streams will happen as the secondary takes over. The *SenderCompID* changes and the *PacketSeqNum* as well as the *MsgSeqNum* for affected products are reset to 1.

3.7.4 Market Data – Enhanced Market Data Interface (EMDI)

A change of Source IP of the affected multicast streams will happen as the secondary takes over. The *SenderCompID* changes and the *PacketSeqNum* in the packet headers is reset to 1. The *MsgSeqNum* remains contiguous. For details how to handle this situation see the Market Data Reference manual.

3.7.5 GUI

There is no visible effect in the Trader or Admin GUI.

3.8 Persistency layer failover

The persistency layer of the T7 trading system is running separately from the matching engine to ensure lowest latency and highest throughput. The Services *TradeManager* (trade broadcast recovery) and *App/Seq* (order status broadcast recovery) are the visible interface parts of the persistency layer and both services are deployed with a standby partner that takes over processing in case the primary component fails.

The failover of these components is completely transparent on all interfaces but sometimes duplicate or missing messages (gaps) might occur on an affected broadcast stream due to the failover. Handling of gaps and duplicates on these streams are already implicitly covered by the ETI application protocol via gap detection plus retransmission and duplicate detection.

Potential message gaps might occur on sequence number based broadcast streams but will not happen for message key based streams.

3.9 Central component failure

The T7 system design caters for an outage of single central component via failing over to a standby process. The complete failure of central component on a trading partition is thus very unlikely.

A central component is the matching engine, EMDI, MDI, EOBI or one of the services from the Persistency Layer (see section 3.8), but not the T7 Entry Service. If any of these components experience a failure, trading on that partition is stopped and the matching engine service becomes unavailable for that partition. This chapter applies also to the failure and re-start of the complete T7 system, i.e. a failure of all central components.

Resolving the situation requires in all cases manual intervention of T7 operations and will take considerably longer than a few seconds as a restart of the matching engine is needed. When the trading services are re-enabled initially all products on the affected partition are in product state *"Halt"*, non-persistent orders and quotes are deleted and trading continues with an auction.

Though the matching engine service will be disabled when a central component fails it depends on the type of failed component in which order (and when) other services on that partition become unavailable between the incident and the re-start. For example, if the *App/SeqStatus* service from the persistency layer fails the *TradeManagerStatus* might still be available and usable for a while until the manual shutdown is triggered. A market reset from the matching engine (see also section 3.5) is also not triggered until the restart, but since the matching engine service is unavailable, no orders can enter the system in the meantime.

Please note that due to the preliminary nature of all responses sent by the matching engine, it is crucial to synchronize the state of all orders with the order restatement in case a market reset has been issued by the matching engine, see 3.5 and 3.9.5 for details.

3.9.1 Matching Engine

A complete matching engine failure manifests itself in a **ServiceAvailabilityBroadcast** message indicating at least the *MatchingEngineStatus* on partition <x> is unavailable. The data below is for partition 1 (FDAX)

```
ServiceAvailabilityBroadcast.PartitionID = 1
ServiceAvailabilityBroadcast.MatchingEngineStatus = 0
```

Depending on the type and order of failure unavailability indications for *App/SeqStatus* and *TradeManagerStatus* are received in the same or additional **ServiceAvailabilityBroadcast** messages.

After a successful restart of all services this will be signaled by another **ServiceAvailabilityBroadcast** message with *MatchingEngineStatus* availability set to 1 for the partition (and when available also for the other services):

```
ServiceAvailabilityBroadcast.PartitionID = 1
ServiceAvailabilityBroadcast.MatchingEngineStatus = 1
```

Trading continues always with a market reset, all products set to product state *"Halt"* and an auction afterwards.

3.9.2 Market Data – Enhanced Orderbook Interface (EOBI)

As EOBI market data is sent out from the matching engine a failure of the EOBI component may only happen if the matching engine fails.

After the re-start, the *AppISeqNum* and *MsgSeqNum* in EOBI messages are both reset to 1 and the field *AppISeqResetIndicator* is set to indicate the restart. The market reset is communicated via an *Order Mass Delete* message. This is then followed by *Order Add* messages of all persistent orders in the book.

3.9.3 Market Data – Enhanced Market Data Interface (MDI, EMDI)

If the EMDI component is the failed central component no further snapshots will be received for products of the affected partition. In case another component failed no new data is received, but snapshots are still received.

There is no impact on MDI other than no new data is received for affected products.

After the re-start the *SenderCompID* will change and the *PacketSeqNum* and *MsgSeqNum* in the packet headers are reset to 1.

3.9.4 GUI

The Service Availability changes are indicated in the GUI status bar and displayed in detail in the Newsboard. Submitting orders for the respective partition during the failure time is not possible and requests will be rejected.

The changed order book on restart is displayed in corresponding market overviews and the product state change to “*Halt*” on restart is indicated additionally in the Newsboard view.

3.9.5 Best Practice

Recovery from a central component failure requires manual intervention by T7 operations and may take a considerable amount of time. The market situation and reopening of the market may include additional explanations and planning from the exchange side. Members are strongly advised to follow the exchanges Newsboard messages via ETI News Broadcast, T7 Trading GUI Newsboard, Newsboard webpages or SMS and Email Emergency Notifications. See section 2 for details on these services.

There will be at least 15 minutes (30 minutes after a 3 hour outage) given to all participants after Order/Quote Management Service of a partition becoming available after a failure of a central component.

This means there will be ample time given before actual trading commences to maintain orders disseminated via the orderbook restatement and review trades via trade broadcast recovery. Orders reported in the restatement are part of the active order book and will become active once trading is resumed. The orders and trades can be (re)inquired via retransmission requests any time after the Order/Quote Management Service and the Trades Service becomes available. We recommend synchronizing the state of all orders with the restatement messages.

We urge participants to synchronize the states of all their orders with the restatement issued after Market Reset has been announced by the matching engine. The reason is that there may have been order modifications, order cancels, order adds, or order executions acknowledged by the matching engine which did not reach the persistency layer and thus are not taken into consideration after a market reset.

Information about legally binding execution of orders or quotes prior to the component failure must only be recovered from Trade Capture Reports on the Trade Broadcast stream. All trades reported via Trade Capture Reports will be unconditionally reported to the Clearing instances of the respective exchange.

Order or quote executions reported only via Execution Reports (or public market data) and not via Trade Capture Reports were not legally executed and according orders/quotes are after a partition restart either deleted (non-persistent orders) or become again part of the active orderbook for successfully persisted orders. Between partition restart and before actual trading is resumed there are no pending “in-flight” Execution Reports or Trade Capture Reports in the T7 system thus there is time to reconcile between these two broadcast streams. See ETI Manual 4.13.3 for technical details how to reconcile between Execution Reports and Trade Capture Reports.

The failure incident on 14th April 2020 generated a unique error situation where a very small amount of executed orders were processed twice by the T7 Trades Services and subsequently reported in TradeCaptureReports with different *TradeIDs(1003)*. These duplicates can be detected by checking for a unique combination of fields in TradeCaptureReports. For Derivatives markets *SideTradeID(1506)* is unique per product and business day, for Cash markets the respective combination *SideTradeID(1506)*, *LeavesQty(151)* and *TradeNumber(2490)* per product and business day. If such a combination is reported twice with different TradeIDs(1003) this indicates a severe problem in T7; please inform Market Supervision about the situation asap in this case to allow for appropriate actions on the exchange side.

3.10 T7 Entry Service failover and failure

When the primary component of the T7 Entry Service fails the standby component takes over in a fully transparent way. The T7 Entry Service broadcasts streams might reveal missing messages (gaps) or messages received twice. The streams are sequence number based so both effects are already implicitly covered by the application protocol via gap detection plus retransmission and duplicate detection.

In case of a complete failure of the T7 Entry Service the according trade- and retransmission services become unavailable. Re-enabling the service requires manual intervention of T7 operations and will take considerably longer than a few seconds. On re-start of the services after a failure all TES products on that partition are initially in TES product state *OFF*.

3.10.1 Order Entry (ETI)

There is no visible impact at the ETI interface in a failover situation except for possibly duplicated or missing broadcast messages.

In case of a failure a **ServiceAvailabilityBroadcast** message indicates the TradeEntryService and its broadcast retransmission service on partition <x> is unavailable. The data below is for partition 1 (FDAX)

```
ServiceAvailabilityBroadcast.PartitionID = 1
ServiceAvailabilityBroadcast.TradeEntryServiceTradeDate = 0
ServiceAvailabilityBroadcast.TradeEntryServiceRtmTradeDate = 0
```

During the service downtime all request submissions to the TradeEntryService or the according TradeEntryServiceRtm (broadcast retransmission service) are rejected.

Once the service is available again an according **ServiceAvailabilityBroadcast** message is distributed with the availability set to 1.

3.10.2 Market Data – Enhanced Orderbook Interface (EOBI)

No visible impact on EOBI.

3.10.3 Market Data – Enhanced Market Data Interface (EMDI)

No impact on EMDI on failover.

On re-start after a failure the initial product states for TES products are reported.

3.10.4 GUI

The Service Availability changes are indicated in the GUI status bar and displayed in detail in the News board. Submitting orders for the failed T7 Entry Service during failure time is not possible and such requests will be rejected.

3.10.5 Best Practice

After a complete failure and restart of the T7 Entry Service applications should check the highest available sequence numbers for all affected broadcast streams and close any potential gaps afterwards via retransmission requests.

3.11 GUI unavailability

When the infrastructure of the T7 GUI is not available client sessions will be potentially logged out and presented with the Login View. Until the infrastructure is available again no login attempts will be successful.

3.11.1 Order Entry (ETI)

There will be no effect on ETI sessions.

3.11.2 Market Data – Enhanced Orderbook Interface (EOBI)

There is no effect on EOBI data dissemination.

3.11.3 Market Data – Enhanced Market Data Interface (EMDI)

There is no effect on EMDI data dissemination.

3.11.4 GUI

Usage of the GUI is not possible, and logins will fail.

3.11.5 Best Practice

Members are advised to follow the exchanges Newsboard messages via ETI News Broadcast, the Newsboard Webpage, or information received via SMS and Email Emergency Notifications. Any news about the expected downtime or availability of the T7 GUI will be reported there. See section 2 for more details on these services.

4. Appendix

4.1 Handling of Matching Engine Service Availability notifications

Some scenarios in this section imply a **ServiceAvailabilityBroadcast** message indicating a change of MatchingEngineStatus (Order/Quote Management Service) for a given partition. As an example, when Order and Quote Management for partition 1 becomes unavailable the following ServiceAvailabilityBroadcast will be sent:

```
ServiceAvailabilityBroadcast.PartitionID = 1
ServiceAvailabilityBroadcast.MatchingEngineStatus = 0
```

This message does not always imply that a matching engine has failed, a failover is in progress or a market reset on that partition is performed. It thus also does not imply that all non-persistent orders and quotes are deleted. The deletion of non-persistent orders and quotes is always reported by a Market Reset (partition scope) or a Mass Cancel Event (session scope) on the session data and listener data streams.

Market Resets occur for example during a matcher failover (see 3.5), automatic Mass Cancels on session scope because of a Low Frequency Gateway failure (see 3.1) or a connection loss between Low Frequency Gateway and a PS Gateway (see 3.2).

Due to race conditions in the T7 system it may happen that the Order/Quote Management Service is shortly unavailable and becomes available again without any other side effect.

4.1.1 Best practice

If the Order/Quote Management Service for a partition becomes unavailable and no Market Reset or Mass Cancel Event is received and deletion of non-persistent orders and quotes is required, the session should send a Mass Cancel Request for the respective required scope.

5. Change log across the Releases

No	Chapter, page	Date	Change
V1.0	General	25-07-2016	Creation
V 2.0	General, 3	01-06-2017	Cash Market migration to T7, Adapted EOBI specific sections as EOBI is disseminated from the matching engine. Added details to EMDI failover (chapter 3.5)
V 3.0	3	10-10-2017	Introduction of Partition Specific Gateways
V 4.0	3	10-09-2018	Removal of Connection Gateways for ETI 7.0
V 4.2	General	18-03-2019	New Version for T7 Rel 7.1 – updated links
V 4.3	General	30-08-2019	New version for T7 Rel 8.0 Descriptions for connection between LF and PS Gateways
V 4.4	General	08-04-2020	New Version for T7 Rel. 8.1 - Links in Chapter 1.3 and Copyright Statement Date updated. No content changes.
V 4.5	3.9.5	15-05-2020	Updated chapter 3.9.5, best practices after central component failure
V 4.6	General	17-07-2020	Clarify persistency of matcher ETI responses
V 4.7	General	13-08-2020	New version for T7 Rel 9.0 – remove information valid only for older releases