Insights into Trading System Dynamics

Deutsche Börse’s T7®

February 2021

Agenda

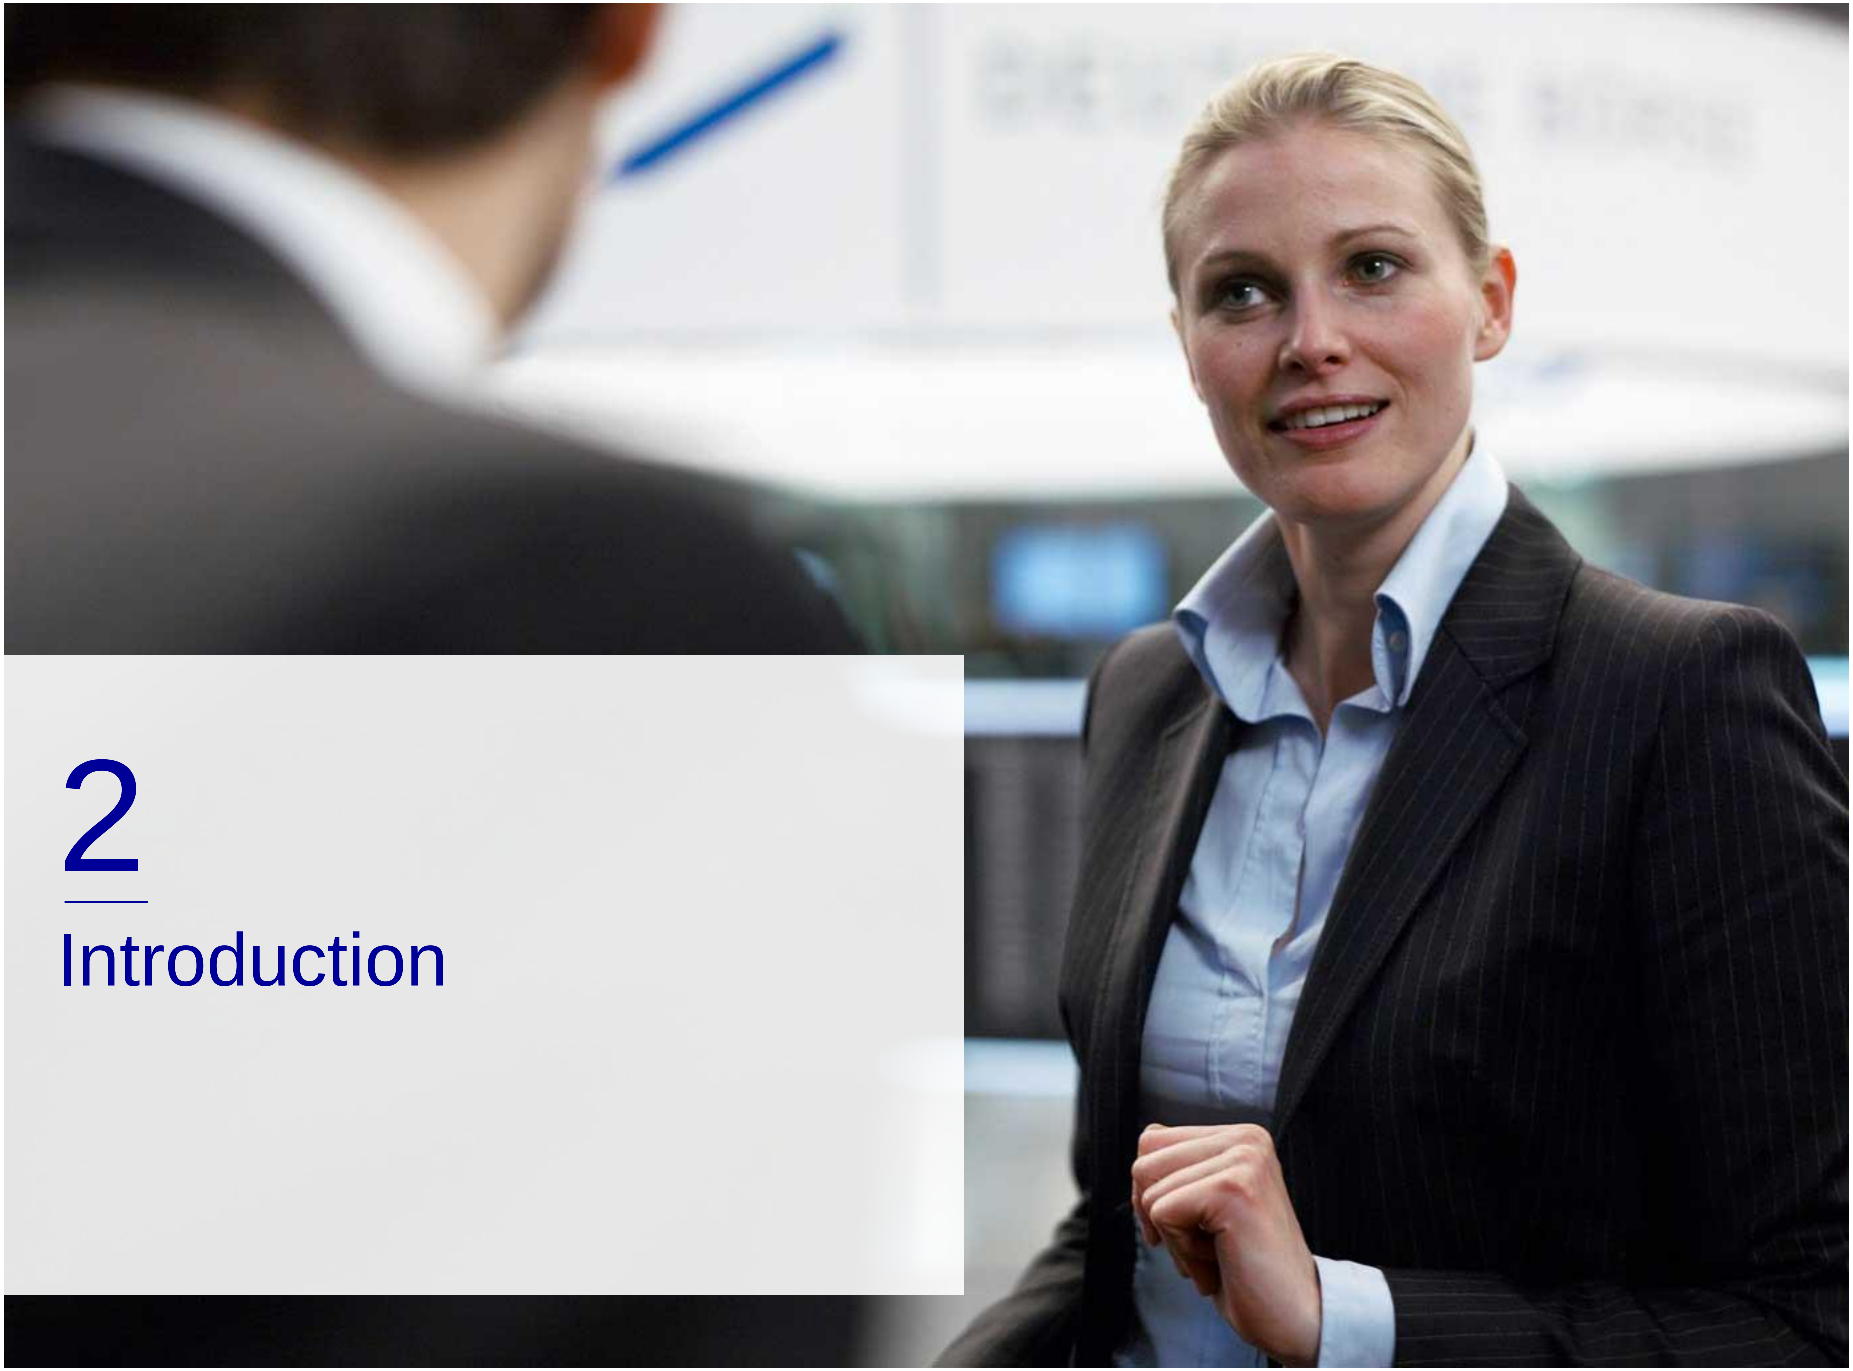

2

Introduction

T7® Technology Roadmap

Deutsche Börse is pursuing its technology roadmap to deliver innovative and superior trading technology.

Recent developments

- Routing of requests entered to Low Frequency (LF) gateways via Partition-specific (PS) gateway enabled for Eurex partition 4 (19 January 2020) and all other Eurex partitions (1 February 2020).
- Release 8.1 introduced on 29 June 2020.
- Measures to curb the negative impact of speculative triggering (DSCP & discard IP) introduced on 13 July 2020.
- Börse Frankfurt/Börse Frankfurt Certificates (XFRA) migrated to T7 on a separate instance (24 August 2020).
- Passive Liquidity Protection introduced on all equity options and DAX® Index Options (24 August 2020).
- Consolidation of all matcher downstream components (Persistency Layer, EMDI publisher, Trademanager) onto single servers in November 2020.
- T7 Release 9.0 introduced on 23 November 2020.
- Consolidation and improved setup for EMDI market data disseminator (7 December 2020).

Outlook

- Q1/Q2 2021: Hardware refresh of Matching Engine servers
- Q2 2021: first rollout steps of optional new FIX gateway architecture
- Q2 2021: Consolidation of PS gateway and Matching Engine process into a single process
- 28 June 2021: T7 Release 9.1

For further details about T7 please visit our websites:
www.eurex.com/t7 and www.xetra.com/xetra-en/technology/t7

T7® Technology Roadmap

Timeline of updates

2020
Jan
Feb
Mar
Apr
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec
2021
Jan
Feb
Mar
Apr
May
Jun
Jul

Activation of LF tunnelling

T7® Release 8.1

EOBI DSCP Flag & discard IP for EUREX

Market XFRA migrated to T7, PLP Introduced for ODAX

T7® Release 9.0
Improved EMDI setup

Start Hardware refresh Matching Engines
First steps of new FIX gateway architecture
Consolidation PS gateway and Matching Engine

T7® Release 9.1

Processed Transactions and Response Times

T7 request – response round-trip times

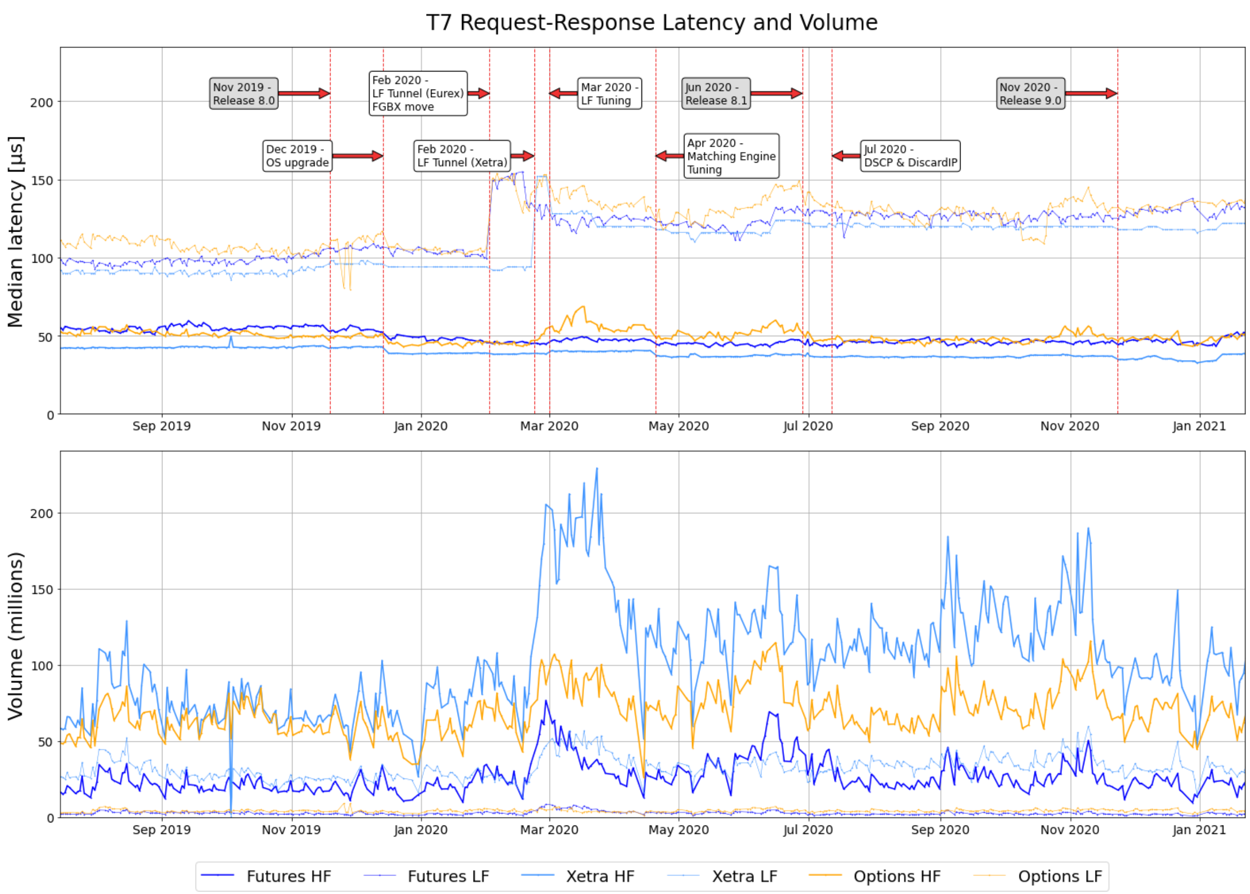

- Deutsche Börse has continuously invested in its trading system and is holding up transparency while providing a low latency trading venue.
- We continued to add functionality while at the same time tuning our system further.
- Move of LF gateways in front of PS gateways led to a temporary increase in roundtrips for LF sessions (February 2020).
- T7 trading system was showing stable performance throughout the exceptionally high volumes in March and April 2020.
- Introduction of discard IP addresses for Eurex reduced number of Matching Engine requests and order book updates.

6

Recent Developments

PS Gateway versus LF Gateways

Latency comparison

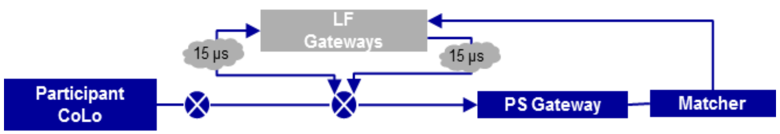

Access Layer (t_3a) to PS gateway in (t_3n)

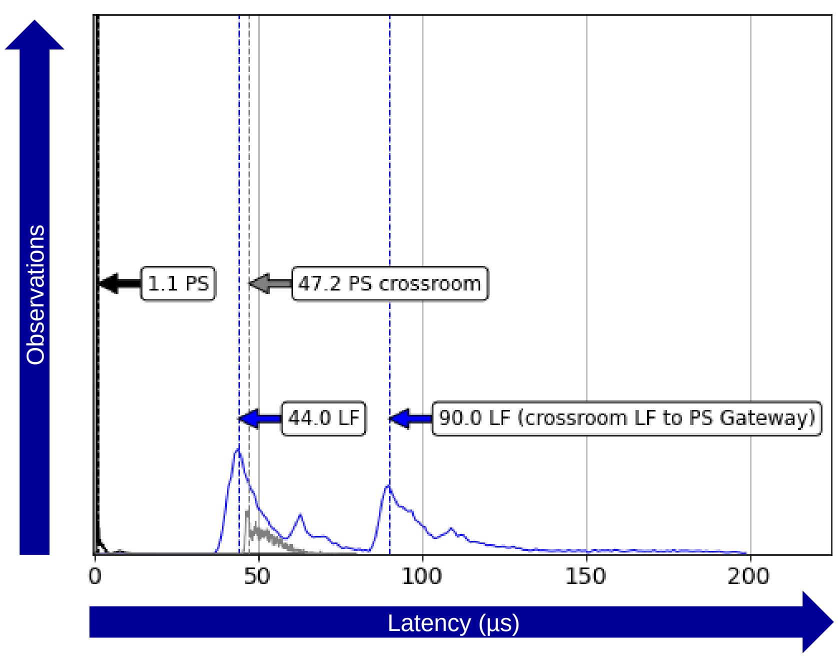

Using 10 Gbit/s cross connects and access via PS gateways provides the fastest way for order and quote management in T7.

LF gateways on the other hand allow access to all partitions of a market via a single session.

Since end of February 2020, all Matching Engine bound requests sent to LF gateways are routed via PS gateways.

The base latency of the path to the PS gateway is around 43 µs higher for LF gateways when compared to directly accessing the PS gateways. Note also that requests that have to cross sides between LF and PS gateways take another ~45 µs longer to reach the Matching Engine.

Consolidation of PS Gateway and Matching Engine

Consolidation Plan

In Q2 2021, the PS gateway and Matching Engine on each partition in T7 will be consolidated into a single process on one server.

The consolidated matcher/gateway processes will deliver faster order processing times and introduce further stability improvements.

The consolidation will enable further reduced latency within the T7 trading platform, simplifies deployment and setup for the exchange and eases failover protocols in case of a component failure.

Introduction

Prior to the architectural change, an upgrade of the server hardware for the T7 Matching Engines will be carried out starting in Q1 2021.

No API changes: The introduction implies no API changes or change of any IP addresses visible to Trading Participants.

Various relative timing figures of order processing and market data dissemination will change during and after the rollout phase.

The detailed introduction plan and further details are outlined in Eurex circular 005/2021 and Xetra circular 004/2021.

New FIX Gateway Architecture

Migration plan to a new FIX gateway architecture

A new FIX gateway architecture will be introduced in T7 starting end of March 2021.

It is based on the existing LF gateway technology and routes requests directly via PS gateways to the respective target partition (so no longer via LF gateways).

Initially, this new FIX interface will be offered for back-office sessions only and in parallel to the existing FIX via LF gateway interface. Trading sessions will follow in a further step.

After a transition period where both FIX interfaces can be used in parallel, the new FIX interface will ultimately replace the existing FIX via LF gateway interface.

The new FIX sessions will be offered only with FIX protocol 4.4.

Introduction Plan

Detailed introduction plans are outlined in Eurex circular 012/2021 and Xetra circular 011/2021.

EOBI DSCP Flags and Discard IP Range

Confining the Effects of Speculative Triggering

The Eurex exchange experienced an increasing amount of "speculative triggered" packets during 2019 and 2020. These packets are sent by Trading Participants based on receiving the first bytes of market data in anticipation of potential trading opportunities.

During processing of further bytes of these market data packets the data of the requests can still be modified to non-trading requests in case the incoming market data turns out not to show a trading opportunity. This is possible if the respective fields of the requests have not yet been written to the underlying physical ethernet layer.

Various adjustments of the Excessive System Usage (ESU) Fees mitigated the problem only for a limited timeframe and the amount of speculative requests became a load issue and an increasing amount of orders went into the orderbook.

The graph shows this development by means of the ratio of sent versus executed IOC orders for Eurex benchmark products.

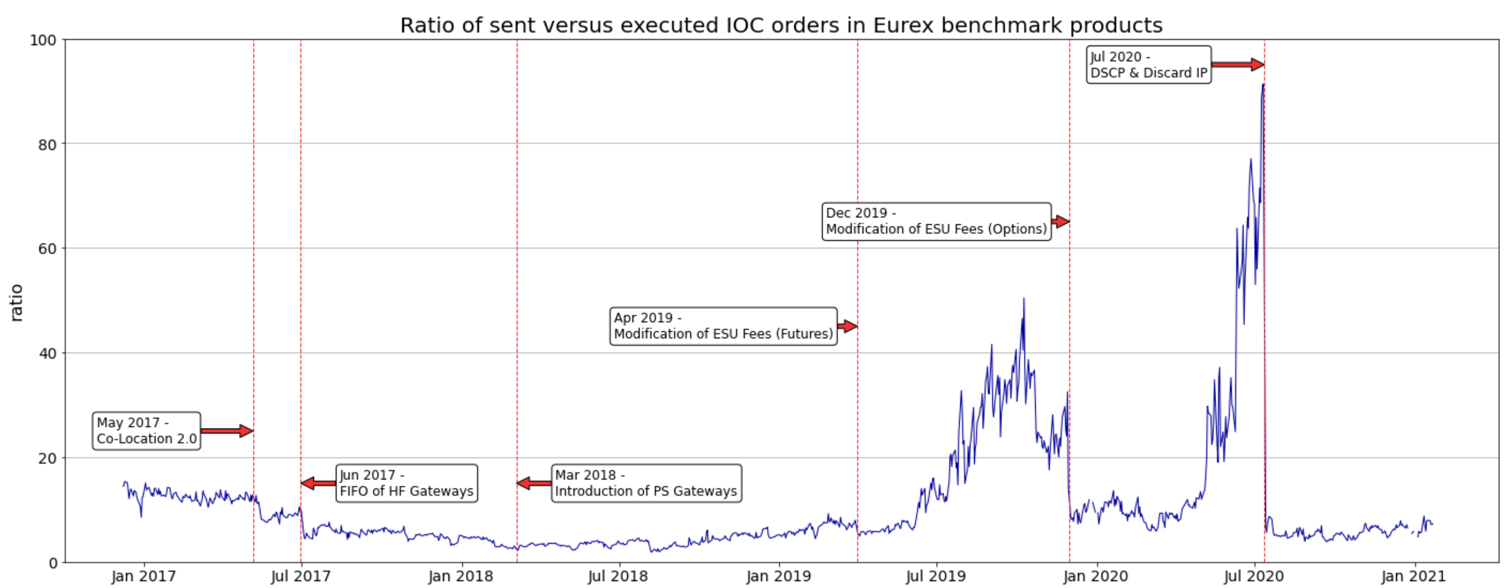

EOBI DSCP Flags and Discard IP Range

Confining the Effects of Speculative Triggering (continued)

Speculative triggering is incentivized by the deterministic technical network architecture of the T7 trading system. The 10 Gbit/s Access Layer Switches operate in a "cut-through" mode and the first bytes of an ethernet frame reserve priority on the competitive network path for the uplink to the next switch. To address this technical problem a technical solution has been implemented.

In July 2020 Eurex introduced a discard IP address range 172.16.0.0/16 on the 10 Gbit/s order entry networks. Trading participants may sent falsely "speculative" triggered packets to the discard IP range, instead of sending it to the exchange. These packets will be discarded at the access layer switch port and no other participants are influenced. Packets sent to the discard IP address are not considered to be orders and are not forwarded to the exchange.

To enable market participants to effectively use this discard IP address, the DSCP field of the IPv4 headers in EOBI market data packets is used. 3 different bits indicate the most common 'interesting' market situations.

The amount packets reaching the Matching Engines decreased significantly after introduction of this measure, while the total number of received packets at the Exchange network boundary did not increase after the introduction of the discard IP address range.

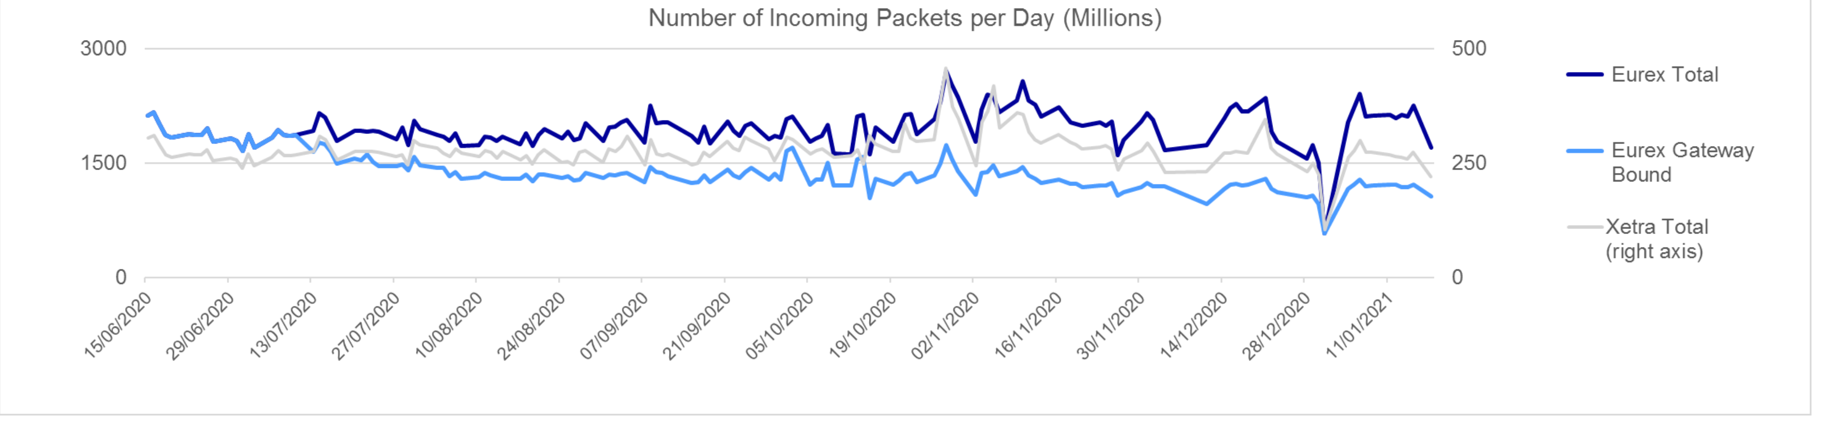

Participant Reaction Time

Insights from the High Precision timestamp file

We define reaction time as the latency between a ‘trigger’ market data packet and a request that leads to an execution.
We use measurement point t_9d for the market data packet and t_3a for the request. We measure 2725 ns as the minimum latency between t_9d and t_3a determined by the T7 infrastructure.

The below charts show the distribution of the reaction times for the FESX to FDAX relation for 27 Sep 2019 (left) and 15 Jan 2021 (right).

We currently measure a minimum reaction time latency of around 2740 ns and a high level of competition (there are around 10 trading participants with reaction times < 2800 ns for most active products).

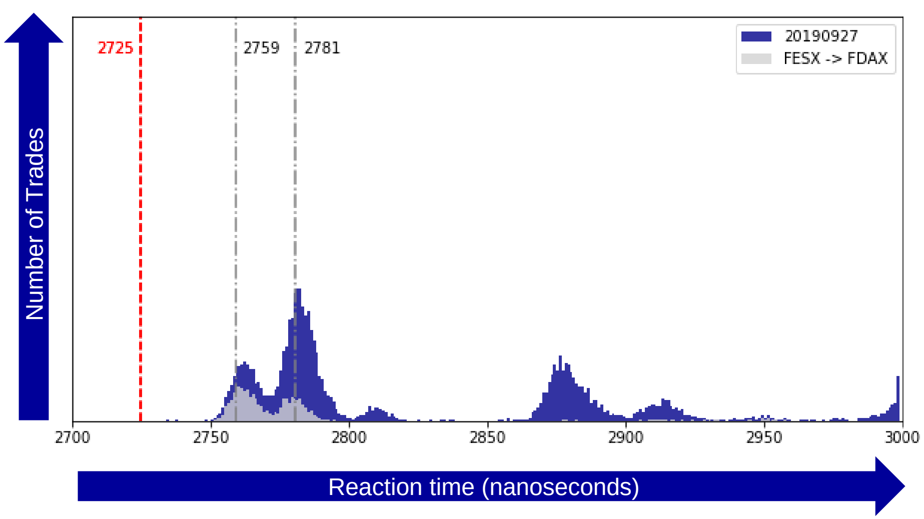

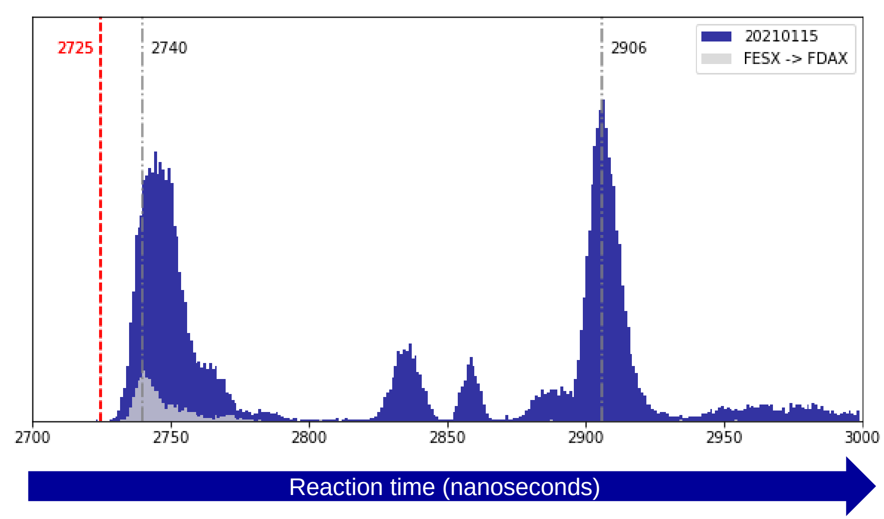

13

Latency Analysis

T7 Topology
At a glance

The below shows the Eurex B side as a schematic example of the topology of the T7 system. Note that Xetra has only two access layer switches.

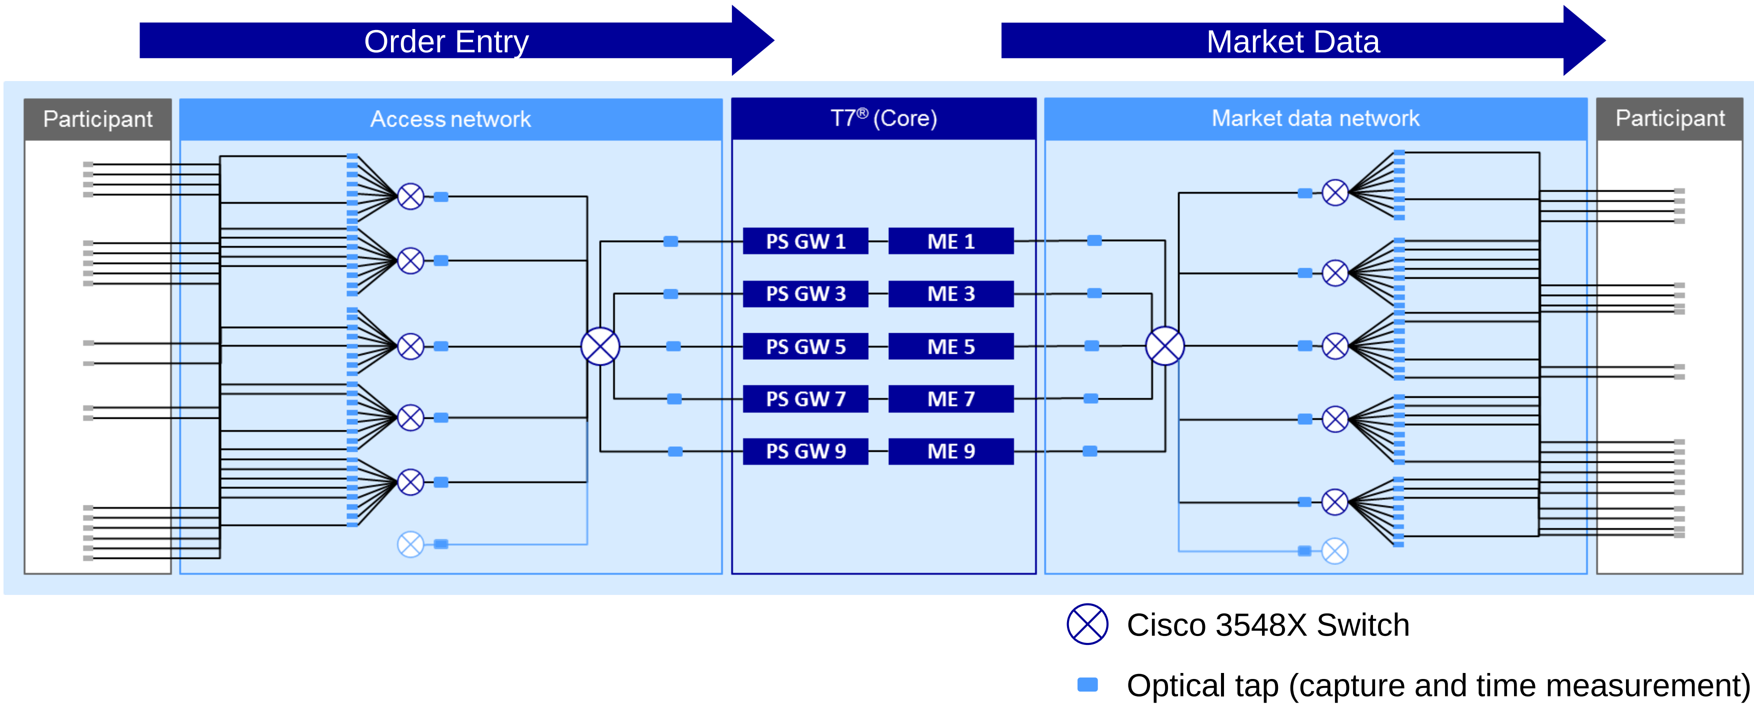

T7 Latency
Composition

The T7 trading system provides utmost transparency about its latency characteristics. Most of the timestamps taken at the various measurement points within T7 Core are included in each ETI response and EOBI market data. With the high precision timestamp file we also make three network timestamps available for each EOBI market data packet (t_3a, t_3d and t_9d).

The latency ‘microclock’ (full circle 60 µs) shows the median latencies for the request-response/EOBI market data path for Eurex futures sent via high frequency sessions measured on 15 Jan 2021.

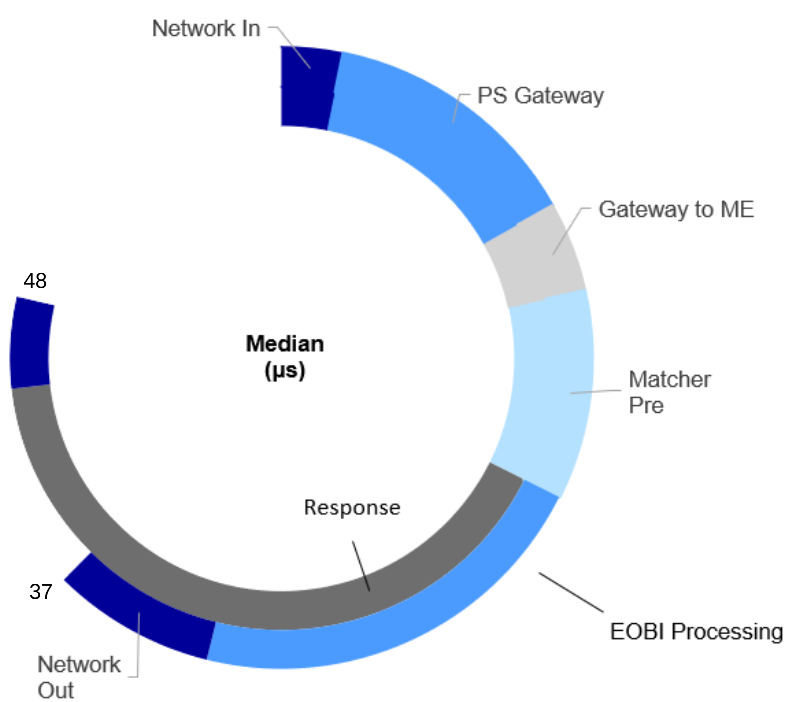

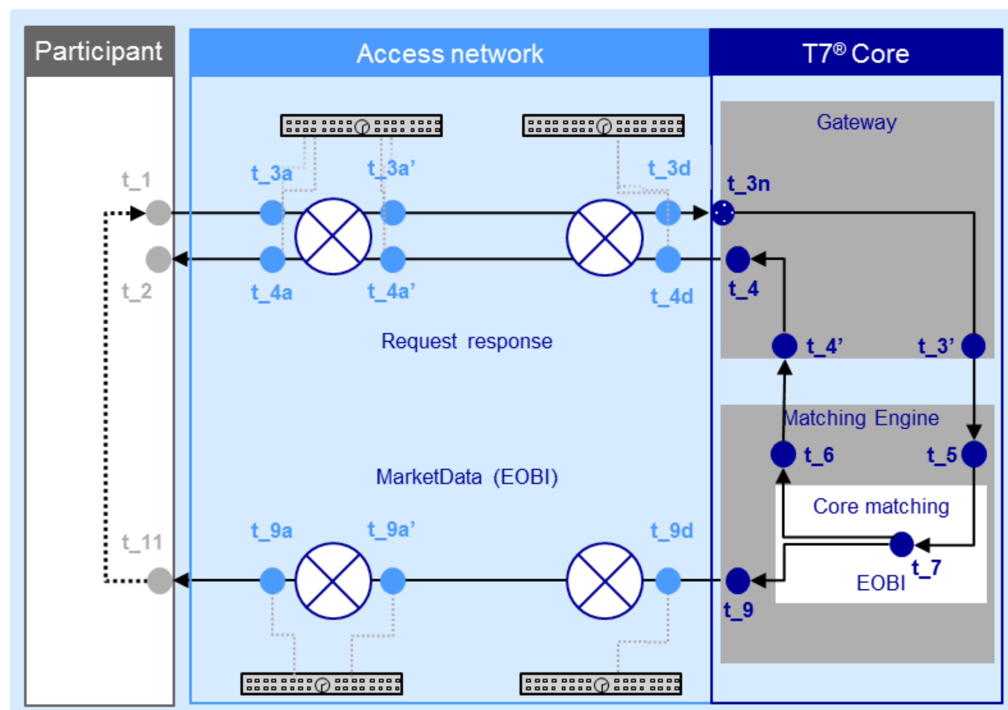

Description of Timestamps

t_[x]a, t_[x]a’	time taken by network capture devices in the access layers.
t_[x]d	time taken by network capture devices in the distribution layers.
t_1, t_2	can be taken by a participant (e.g. via a network capture) when a request/ response is read from/written to the network.
t_3n	time taken by the PS gateway when the first bit of a request arrives on the PS gateway NIC; contained in (private) ETI response for PS Gateway enabled partitions. Consecutive messages via the same session may be assigned to the same t_3n in rare cases.
t_3	time taken by the ETI gateway application when a request is read from the socket on the Participant´s side of the gateway; contained in (private) ETI response for transactions for non PS gateway enabled partitions (e.g. XVIE).
t_4’	time taken by the ETI gateway when a response/ notification is received by the ETI gateway from the matching engine; contained in (private) ETI response/ notification.
t_4	time taken by the ETI gateway when a response/ notification is written to the socket on the Participant´s side of the gateway; contained in (private) ETI response/ notification.
t_5, t_6	time taken by the matching engine when a request/response is read/written; contained in (private) ETI response.
t_7	time at which the matching engine starts maintaining the order book
t_8	time taken by EMDI publisher just before the first respective UDP datagram is written to the UDP socket.
t_9	time taken by EOBI publisher just before the first respective UDP datagram is written to the UDP socket.
t_10, t_11	can be taken by a Participant (e.g. via a network capture) when a UDP datagram is read from the UDP socket.

T7 Latency

Composition (continued)

The below charts show a comparison of current latencies with the spectrum from the last update (August 2019). The data is for Eurex futures sent via PS gateways.

‘Network response/EOBI’ include the TCP/UDP stack on the respective server.

Note that a change of DSCP flag will lead to a ~2 µs higher latencies for some EOBI packets.

Request path and market data (EOBI)

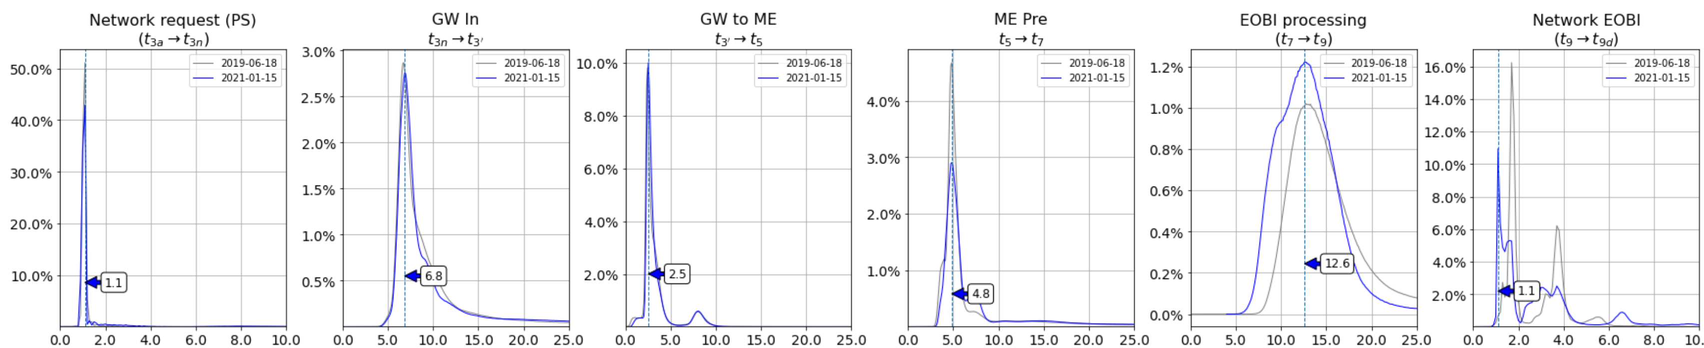

Response path

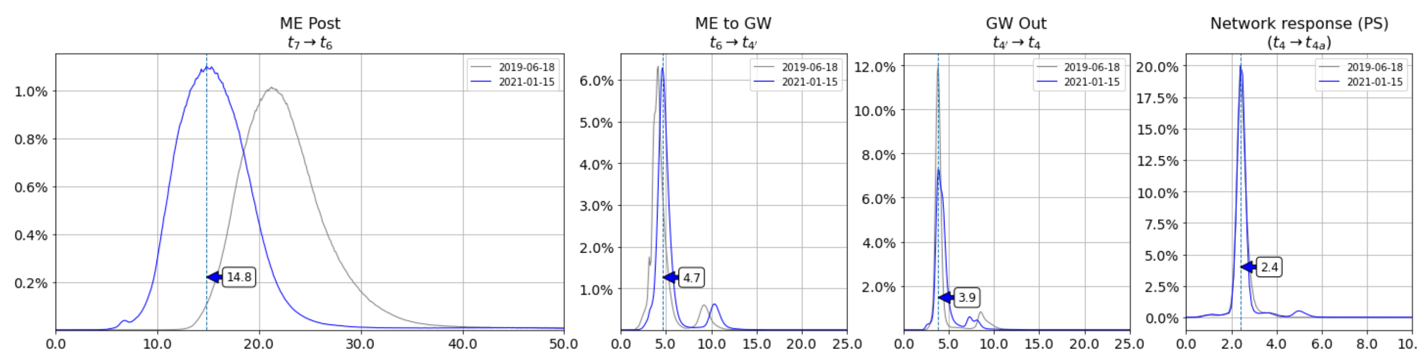

T7 Latency

Microburst dynamics

Microburst dynamics FESX (15 Jan 2021)

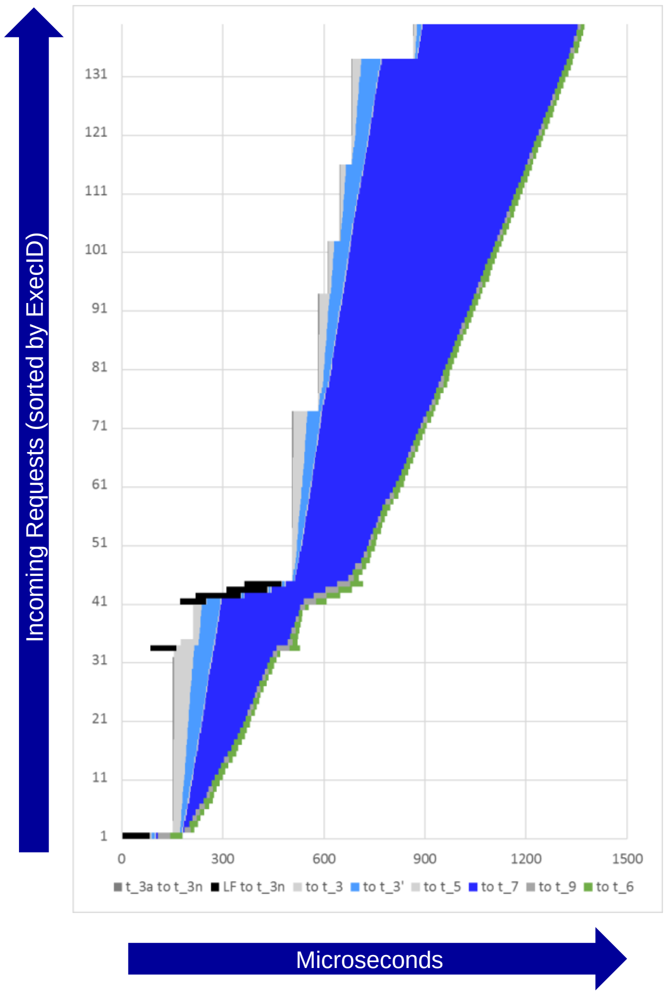

During micro-bursts, the input into the trading system may be greater than the throughput capabilities. This in turn causes queuing which results in higher latencies.

Higher latencies cause risk (i.e. it takes longer to place/pull an order).

T7 provides real-time performance insights by providing up to seven timestamps with each response and key timestamps with every market data update.

The chart on the left shows the paths:

- Network access layer (t_3a) or LF gateway In (t_3) to PS gateway In (t_3n) to
- Matching engine in (t_5) to
- Start matching (t_7) to
- EOBI SendingTime (t_9) [where available] to
- ETI SendingTime (t_4).

Typical throughput rates in kHz (1/ms) are 8000 at t_3n, ~400 at t_3' and t_5, 150 at t_7 and 150 to 220 at t_4.

EOBI send times are usually well before the gateway send time of responses.

Note that base latency for requests entered via LF gateways is ~ 43 µs higher. As all requests are now routed via PS gateways no overtaking between LF gateway and PS gateway requests is observed any more.

T7 Latency

Order Entry Network

The order entry network in Colocation 2.0 uses a two staged hierarchical funnel in approach.

All cables are normalized to guarantee a maximum deviation of +/- 0.5 m (+/- 2.5 ns) between any two cross connects to the exchange.

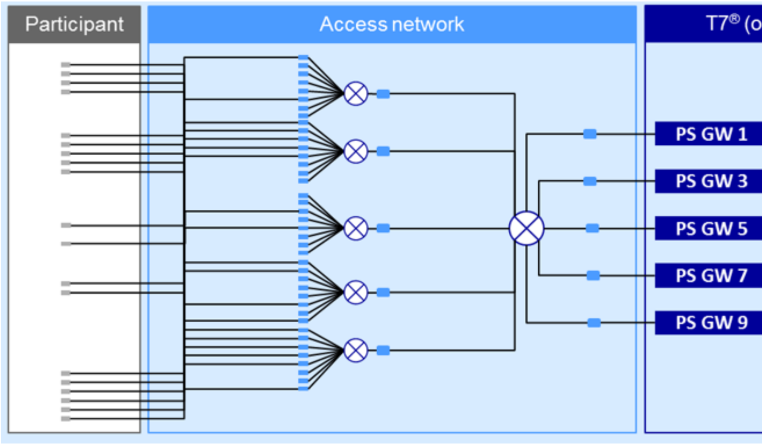

Every inbound and outbound packet on this path is captured via passive network TAPs at 3 different locations.

The packets are timestamped with nanosecond resolution and accuracy.

This capture infrastructure allows early detection and analysis of many kinds of technical network problems and an in-depth latency analysis on network level, like overtaking probabilities between measurement points.

The high precision timestamp file service enables participants access to timestamps t_3a, t_3d and t_9d for each request leading to a an EOBI market data update.

Network Access
Order entry latency

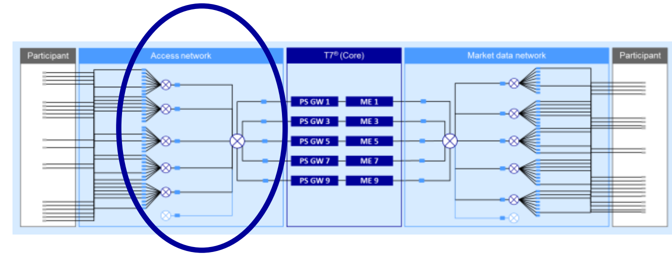

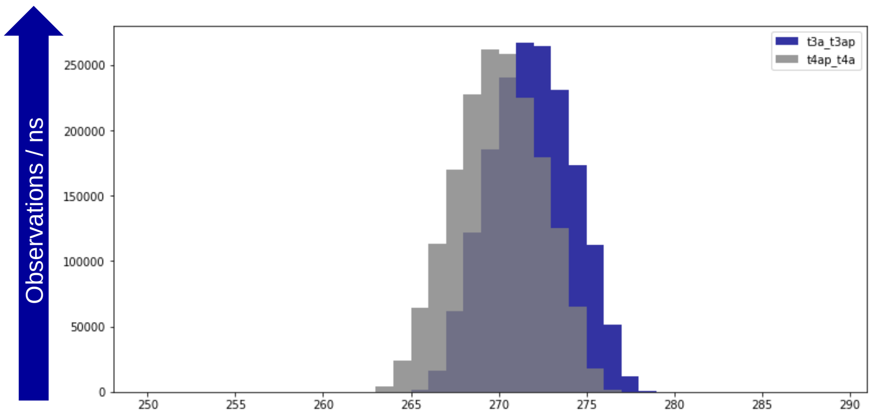

The base latency and latency jitter is identical for all access layer switches within the measurement precision. The spectrum is very tight with a variance of under 3 ns (top left and inset).

Note that bursts of traffic may lead to queues in access and distribution layer switch because multiple lines funnel into a single uplink (bottom left, note the logarithmic scale).

The chart on the bottom right shows the overtaking probability on a switch between a first message in a burst and consecutive messages. The overtaking drops sharply and there is no observed overtaking beyond 3 nanoseconds delta between messages.

Data is taken from the Eurex B side.

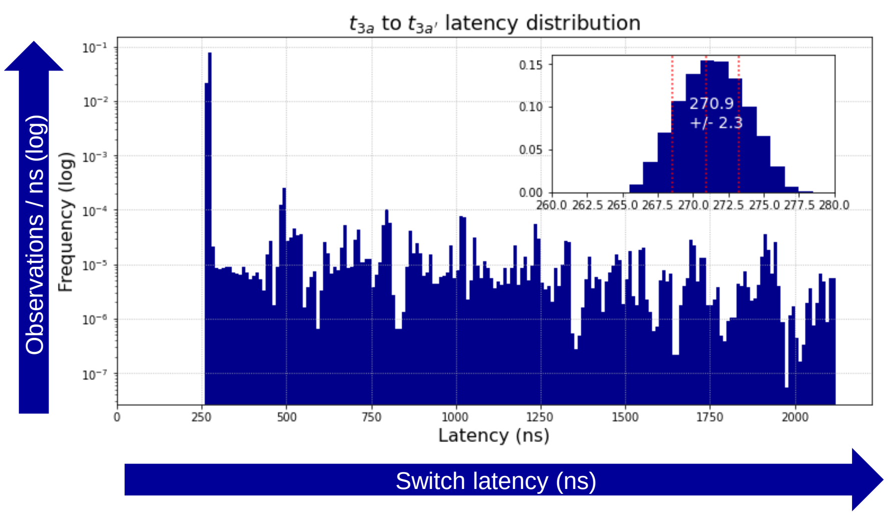

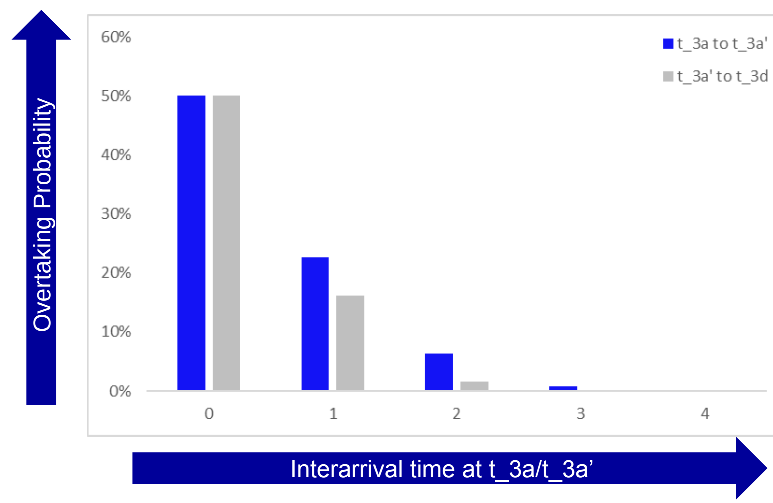

Network Access

Market data distribution

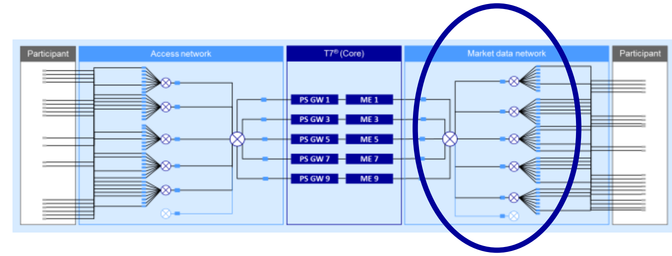

The market data network has a funnel in – fan out topology. It funnels in data from different market data disseminators (on the distribution layer switch) and fans it out via multiple access layer switches.
We took extra care to ensure a fair and deterministic network.

The data volume flowing from the distribution layer to each access layer switch is identical. The switch model was chosen for its deterministic latency profile (t_9d to t_9a distribution is ≈15-20 ns) and for its minimal port to port jitter.

We currently observe a semi static latency difference of up to 10 ns (comparing t_9a of two different ports on the same switch) due to internal processing of multicast by the switch.
The Cisco 3548X switch divides the on board Ethernet-Ports into three buffer areas. Serial replication for multicast takes place within these areas. In terms of latency this means any one port within this area could be faster than other ports, which is an artifact of variable internal pointers to packet queues. The starting port for replication of a packet in a buffer area is identified by an internal pointer.

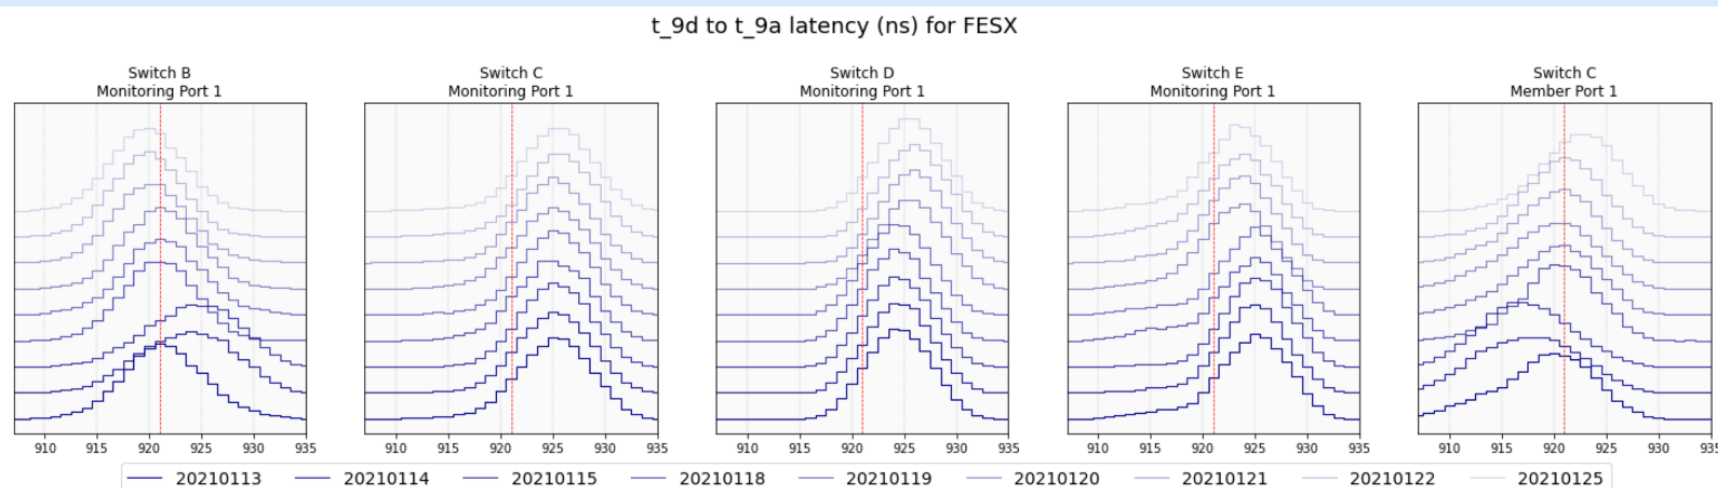

Latency

Comparison of access types

The table below gives an overview of current latency figures of the T7 trading system. All times given are in microseconds.

Network timestamps (t_3a, t_4a) are synchronized using PPS and white rabbit. The time synch quality between these timestamps is thus ~1ns. Other T7 timestamps are subject to jitter of up to +/- 50 ns.

The difference in median latency of the inbound path (t_3a to t_3n) of LF vs PS gateways is 73 µs for Eurex and 38 µs for Xetra.

Market	Gateway Type	percentile	t_3a to t_3n	t_3n to t_5	t_5 to t_4	t_4 to t_4a	t_3a to t_4a
Eurex	PS	50%	1.0	11	33	2.4	48
	LF	50%	74	11	69	11	171
	Delta LF - PS	50%	73	< 1	36	8.6	123
Xetra	PS	50%	1.0	11	35	2.4	49
	LF	50%	39	15	67	11	138
	Delta LF - PS	50%	38	4	32	8.6	89

Eurex	PS	25-75%	<0.1	5	26	<0.3	34
	LF	25-75%	51	9.2	22	1.2	60
Xetra	PS	25-75%	<0.1	5	34	0.3	62
	LF	25-75%	27	25	27	0.9	51

Orders/quotes – detailed performance data

Our transparency

For the top 15 futures products, daily statistics about the matching engine processing times as well as gateway processing times are provided via the 'Member Section' on Eurex Exchange's website. The ETI round-trip times are calculated based on t_4 – t_3n (gateway SendingTime – gateway RequestIn).

The table below additionally contains latency figures for DAX equities. All data displayed below refers to 15 January 2021.

Product Description	Symbol	Matching engine Round-trip times (in µs)			Enhanced Trading Interface Round-trip times (all GWs, t_4 - t_3 in µs)			Enhanced Trading Interface Round-trip times (PS GWs, t_4 - t_3n in µs)		
		Average	Median	99th percent	Average	Median	99th percent	Average	Median	99th percent
EURO STOXX 50® Index Futures	FESX	100	28	1000	1019	147	62	1000	950	146
STOXX® Europe 50 Index Futures	FSTX	126	24	2500	17	193	56	3250	12	215
DAX® Futures	FDAX	56	26	382	839	98	60	476	767	92
Mini-DAX® Futures	FDXM	57	27	404	703	99	60	464	638	92
MDAX® Futures	F2MX	94	20	2500	34	148	45	3000	28	148
SMI® Futures	FSMI	32	23	176	240	59	48	232	231	56
Euro-Bund Futures	FGBL	110	37	1000	691	154	86	1000	635	154
Euro-Bobl Futures	FGBM	82	47	750	351	121	88	750	336	120
Euro-Schatz Futures	FGBS	46	23	472	69	89	54	750	53	78
Euro-Buxl® Futures	FGBX	96	48	750	475	137	88	750	454	137
Long-Term Euro-BTP Futures	FBTP	58	31	344	207	99	66	430	185	92
Euro-OAT Futures	FOAT	58	30	382	407	95	66	446	379	91
EURO STOXX® Banks Futures	FESB	82	23	1500	121	119	45	1750	115	118
VSTOXX® Futures	FVS	105	27	1500	143	161	60	2000	111	165
STOXX® Europe 600 Index Futures	FXXP	140	26	3250	254	190	54	4000	243	192
DAX® Equities		64	30	750	4310	94	56	750	3886	88

24

Market data

Trading system dynamics

Latency characteristics of EOBI versus ETI - Futures

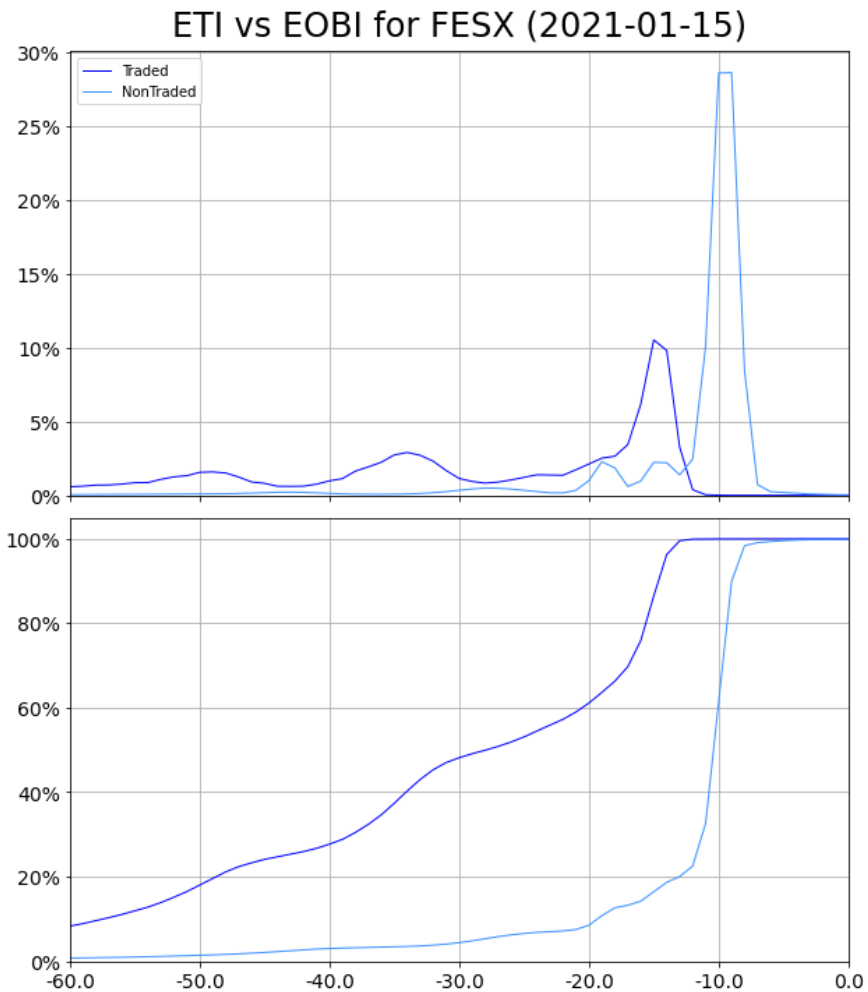

T7 is designed to publish order book updates first on its public data feed.

The diagram shows the time difference distribution between public and private data in microseconds (EOBI first datagram vs ETI responses, t_9 - t_4).

The data is a production sample from 15 January 2021.

EOBI market data is around 10 µs faster than the ETI response for order book updates and more than 15 µs faster for executions.

The first EOBI datagram is faster in more than 99.9 percent of the cases compared to the ETI response and also the first passive ETI book order notification for simple transactions.

The “public data first” principle will be ensured also with the introduction of the consolidated matcher/gateway processes. Details in timing might change after the introduction.

Trading system dynamics

Latency characteristics of EOBI versus ETI - Options

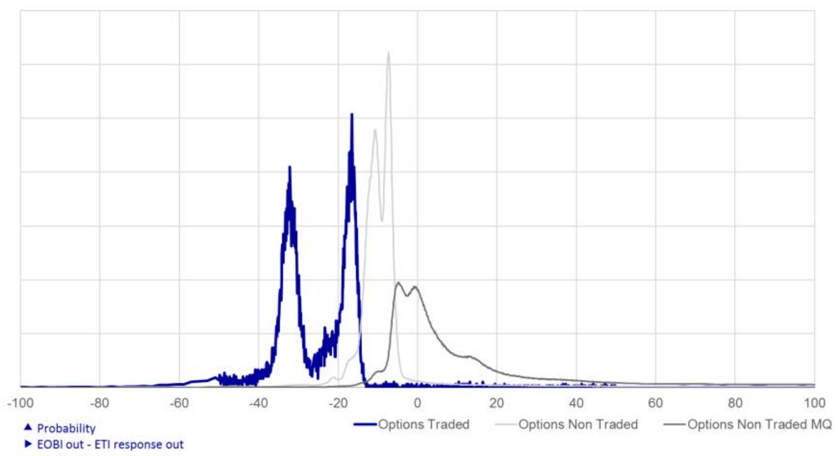

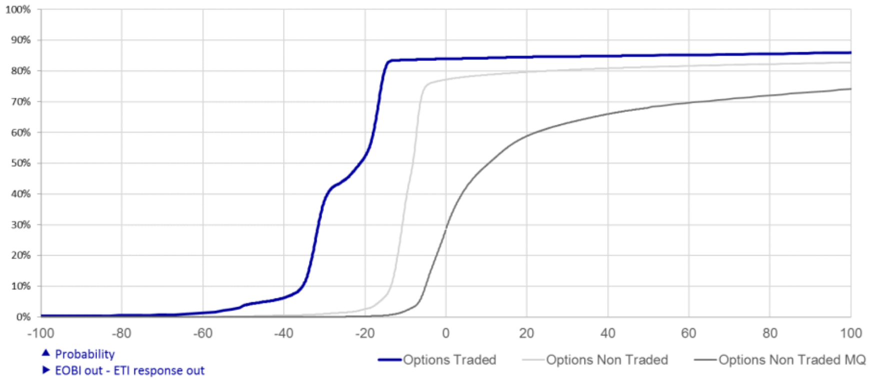

The data is a production sample from 15 January 2021 for OESX Options.

We distinguish between orders leading to a trade, mass quotes with more than one quote pair and orders and single quote updates.

Trades and single updates to the orderbook are received first on EOBI in around 80% of the cases with a median latency advantage of 8 µs (orders) and 20 µs (trades).

There are two main reason for EOBI delays:

1) The transaction is delayed by preceding messages (queues)
2) The transaction causes a market maker protection with many quote deletions

The latency profile for mass quotes is dominated by larger mass quotes, where the EOBI publisher has to broadcast each quote, leading to longer delays and queues in the EOBI path, while the ETI path only deals with a simple mass quote ack.

Trading system dynamics
Latency characteristics of EOBI versus ETI - Xetra

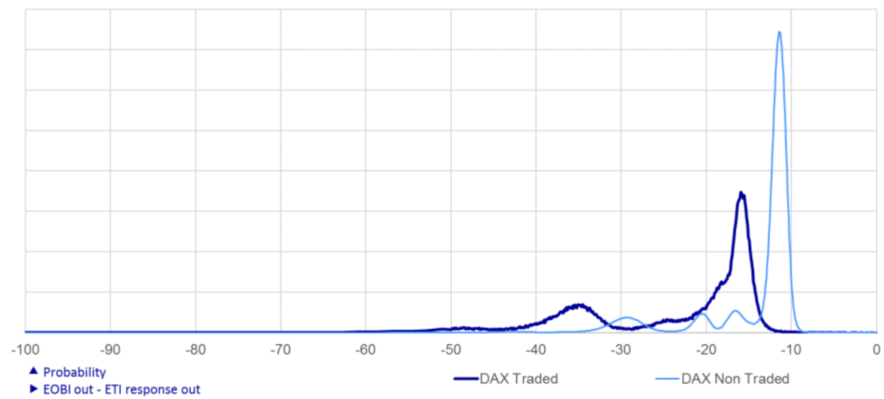

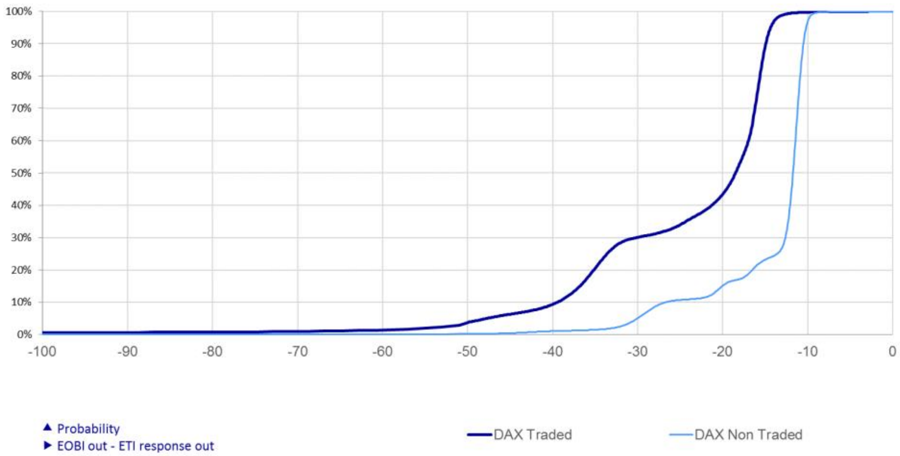

The diagram shows the time difference distribution between public and private data in microseconds for XETRA DAX products (EOBI first datagram vs ETI responses, t_9 - t_4).

The data is a production sample from 15 January 2021.

The latency distribution is very similar to the Eurex futures, simple orderbook updates are received on EOBI 12 µs faster, whereas trades are usually 18 µs faster on EOBI.

Trading system dynamics

Latency characteristics of EOBI versus EMDI

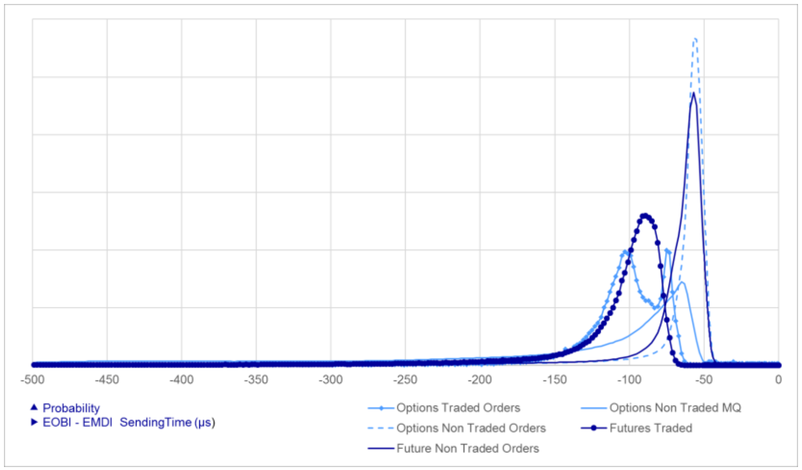

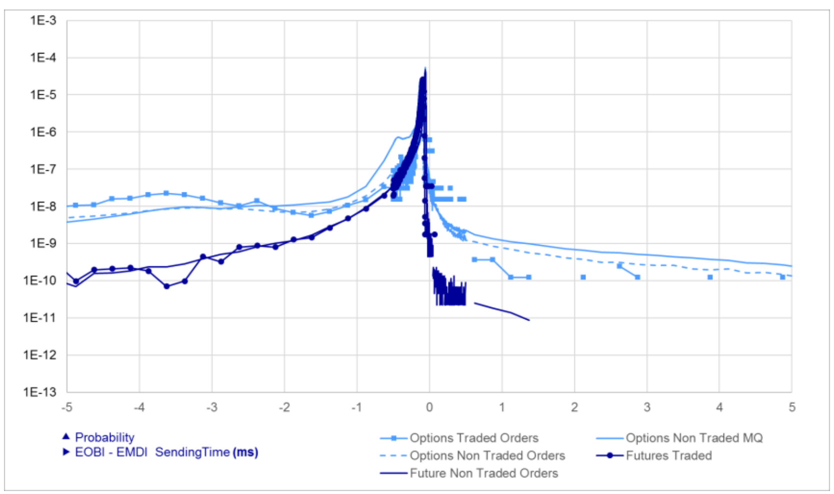

Market data updates provided via EOBI are almost always faster than EMDI updates.

The diagrams show the distribution of t_9 minus t_8, i.e. EOBI first datagram versus EMDI sending time for 15 January 2021 (negative values: EOBI is faster).

The lower diagram is a log plot of the same distribution (the x-axis scale is in milliseconds). It shows that EMDI is significantly slower especially for mass quotes.
At the same time there are some events where option market data is slower on EOBI. The reason is that EOBI slows down when there is a massive amount of individual updates. The table below shows the delay at various percentiles in µs.

Percentile	Options Traded Orders	Options Non Traded MQ	Options Non Traded Orders	Futures Traded	Future Non Traded Orders
0.1	-8125	-7875	-7875	-2375	-4125
1	-5375	-3875	-4375	-499	-499
10	-231	-485	-187	-155	-89
25	-119	-265	-65	-113	-81
50	-101	-111	-57	-95	-73
75	-83	-73	-51	-85	-67
90	-71	-61	-50	-79	-63
99	-63	-40	-38	-71	-57
99.9	155	13125	6125	-67	-53
EOBI faster	99.8%	99.2%	99.4%	100.0%	100.0%

Trading system dynamics

Latency characteristics of ETI versus EOBI versus EMDI

This diagram displays the dependency of the median latency on the complexity of a trade for ETI (t_4 - t_7), EMDI (t_8 - t_7) and EOBI (t_9 - t_7). Note that for ETI we display the gateway sending time of the first passive notification and for EOBI the sending time of the UDP datagram containing the Execution Summary message.
In about 99.99% of all trades, we disseminate order book data on EOBI first (also for larger trades).
Consolidation and improved setup of the EMDI disseminator on 7 December 2020 reduced the EMDI latency depending on the partition by 15 to 30 µs.

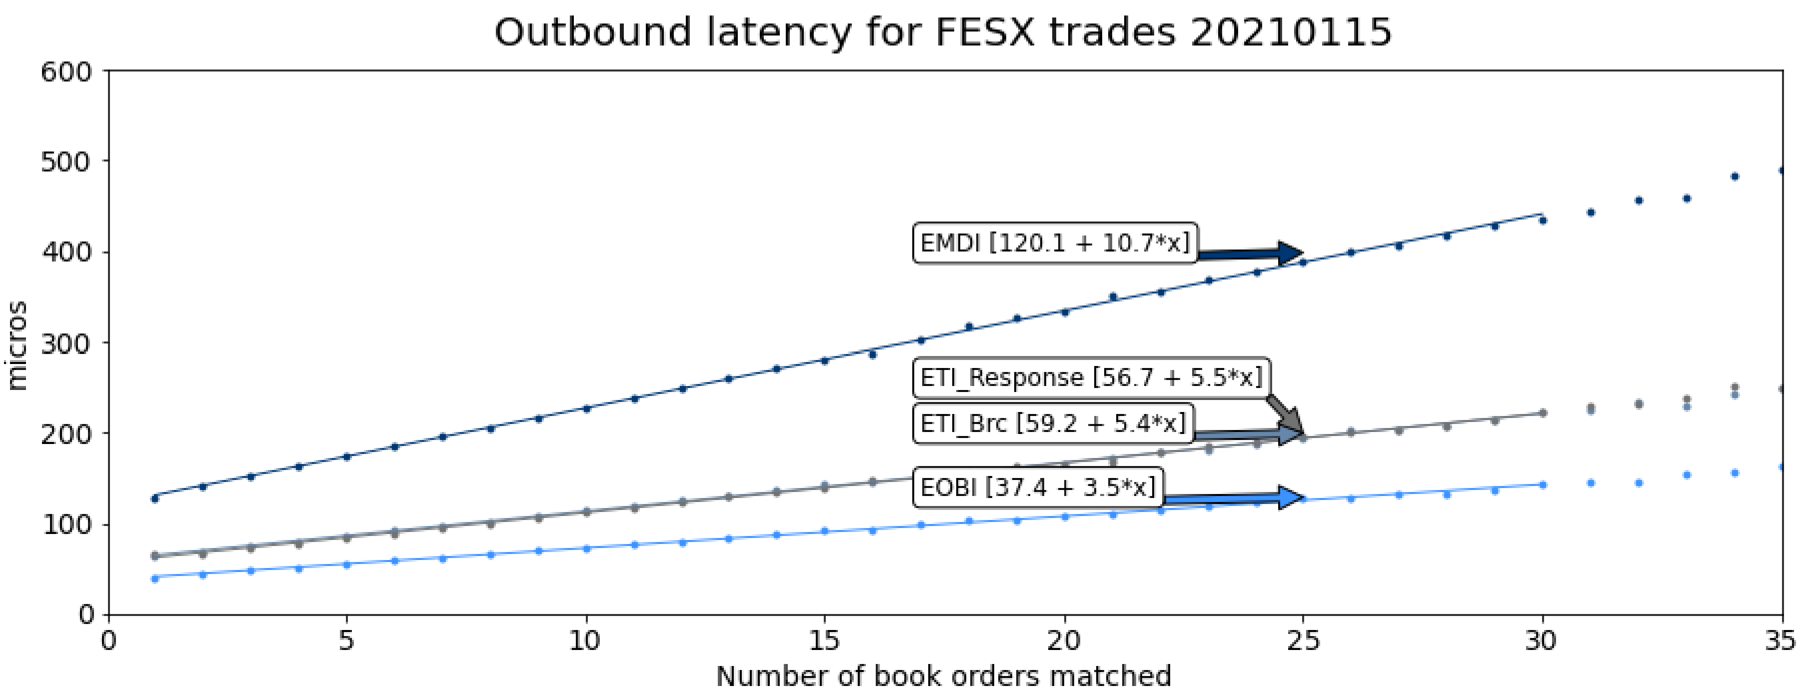

Market data distribution

EOBI latency difference of primary and secondary feed

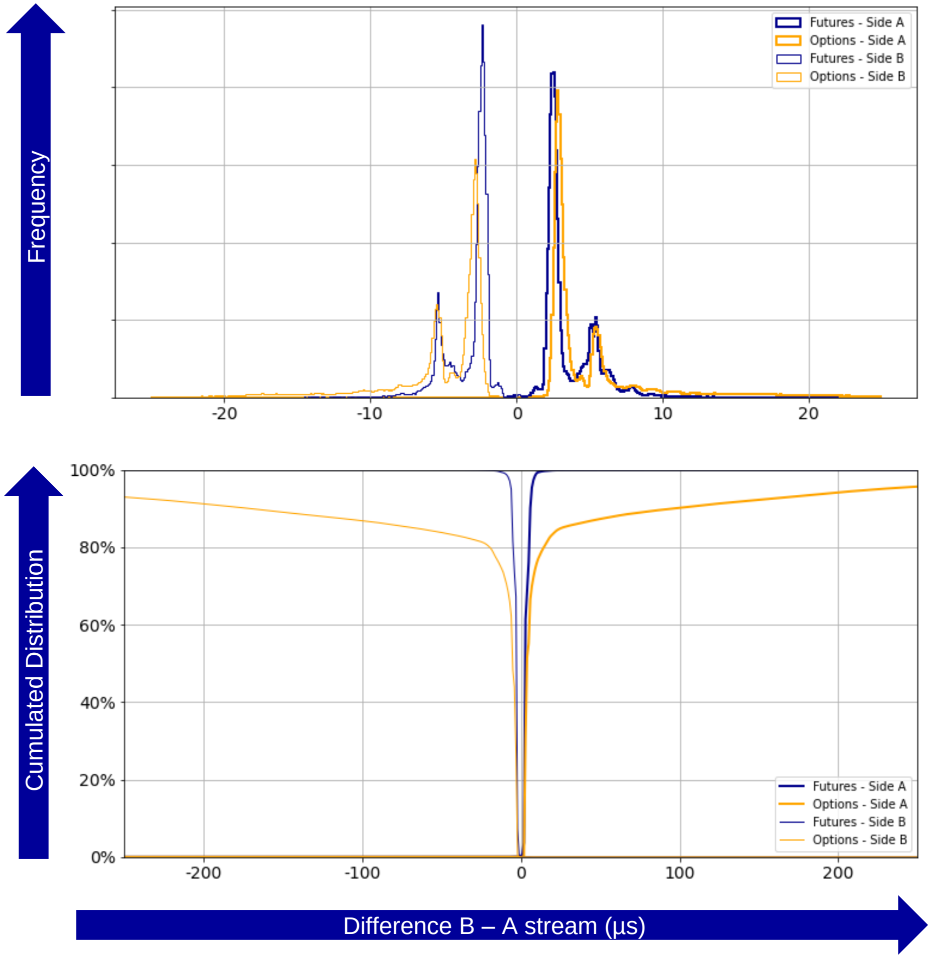

For products assigned to an even partition, market data is published first on the A stream and then on the B stream whereas, for products assigned to an odd partition market data is published first on the B stream and then on the A stream.
The partition ID / product ID is contained in the UDP datagram header of the order book incremental messages and can be used for filtering on UDP datagram level for EMDI / EOBI.
Furthermore, a UDP datagram on the T7 EMDI / EOBI order book delta or snapshot channel will only contain data of exactly one product (e.g. EURO STOXX 50® Index Futures).

The data for both primary and secondary streams is disseminated by the same server using two ports connected to the two sides of the network. The process sends the data first on the primary interface. After all datagrams of a transaction are sent it starts sending on the secondary interface.

The median latency difference between the A and the B EOBI incremental feed is about 2 µs for futures and cash products. For options the median is slightly higher and there are far more outliers (i.e. much slower secondary feed) as the datagrams are published on the secondary feed only after all datagrams of a transaction are sent on the primary feed. As a single transaction may contain many datagrams this time varies by complexity of the transaction, i.e. a mass quote with 200 single quote updates will lead to a higher delay than a single order entry.

Market data distribution

EMDI latency difference of primary and secondary feed

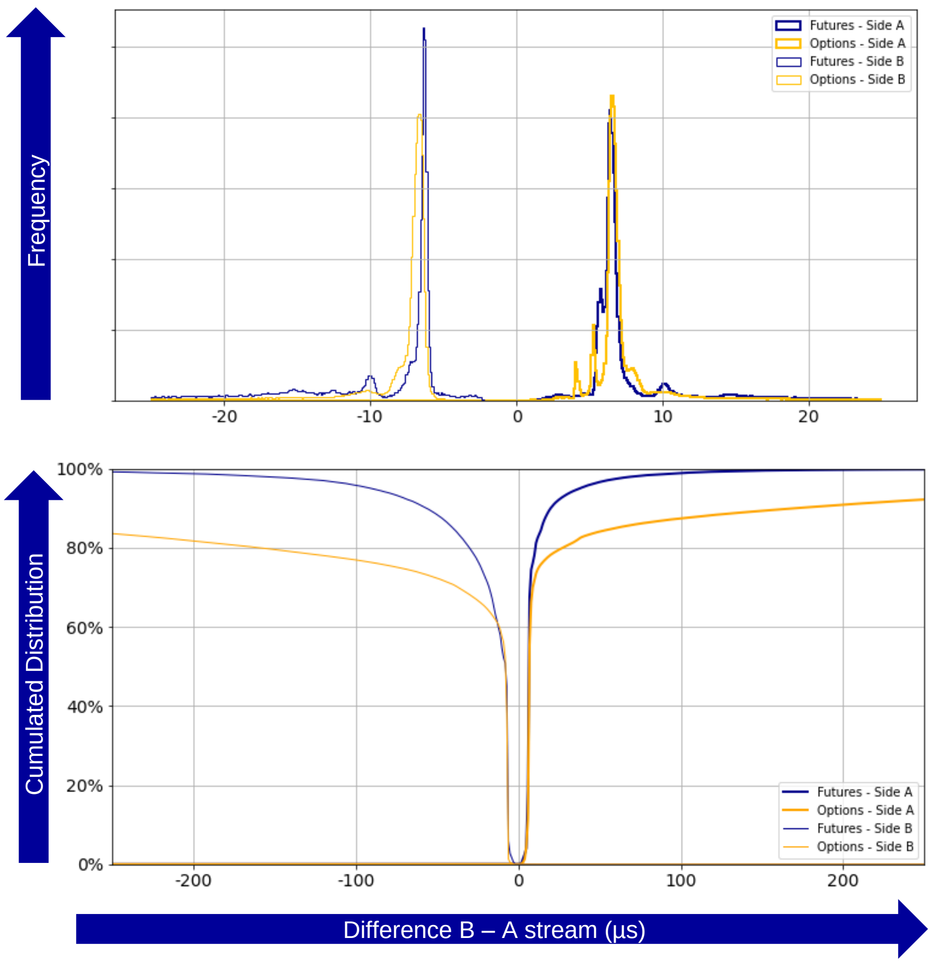

Please note that for products assigned to an even partition, market data is published first on the A stream and then on the B stream whereas, for products assigned to an odd partition market data is published first on the B stream and then on the A stream.
The partition ID / product ID is contained in the UDP datagram header of the order book incremental messages and can be used for filtering on UDP datagram level for EMDI / EOBI.
Furthermore, a UDP datagram on the T7 EMDI / EOBI order book delta or snapshot channel contains data of exactly one product (e.g. EURO STOXX 50® Index Futures).

The data for both primary and secondary streams is disseminated by the same server using two ports connected to the two sides of the network. The process sends the data first on the primary interface. After all datagrams of a transaction are sent it starts sending on the secondary interface.

The median latency difference between the A and the B EMDI incremental feed is about 6 to 8 µs for futures and cash products. For options the median is slightly higher and there are far more outliers (i.e. much slower secondary feed) as the datagrams are published on the secondary feed only after all datagrams of a transaction are sent on the primary feed. As a single transaction may contain many datagrams this time varies by complexity of the transaction, i.e. a mass quote with 200 single quote updates will lead to a higher delay than a single order entry.

Eurex: Market data volume

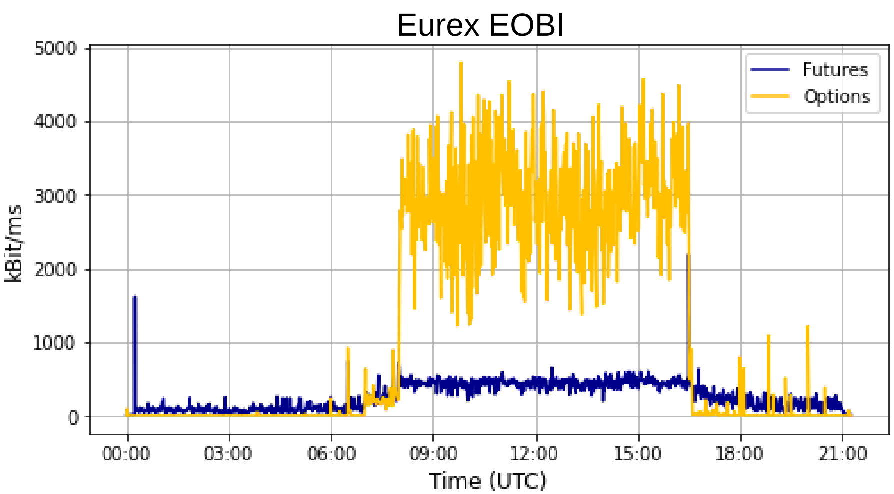

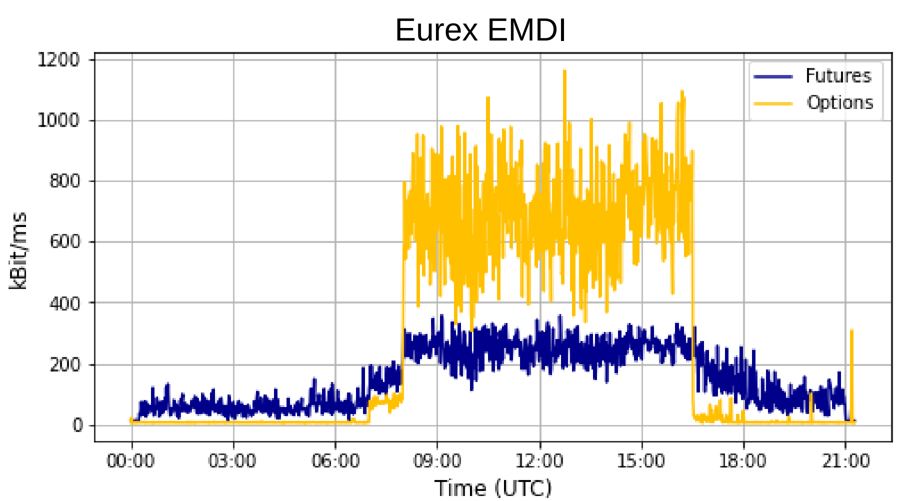

Each data point equals the maximum bandwidth produced on a 1 millisecond scale by the incremental B stream in Mbit/ms.

The provided data shows one data point per minute for 15 January 2021.

Enhanced Order Book Interface (EOBI) peak volume is significantly higher than price level aggregated data volume EMDI. EOBI market data is therefore currently only available to Trading Participants using 10 Gbit/s connections.

The EOBI for options incremental data stream peaks around 5 Gbit/s on millisecond level, while the futures stream peaks at 1 Gbit/s outside market opening.

Participants that want to receive data for Eurex Exchange's products on EMDI with less than 1 ms queuing delays need to use a connection with a bandwidth of more than 1 Gbit/s (options) or 400 Mbit/s (futures) respectively.

Trading Participants are advised to use two 10 Gbit/s connections (one for each market data stream) in Co-Location to receive market data.

Xetra: Market data volume

Each data point equals the maximum bandwidth produced on a 1 millisecond scale by the incremental B stream in Mbit/ms.

The provided data shows one data point per minute for 15 Jan 2021.

Market Data peak data volume can be a significantly higher on high volume trading days. Hence Participants that want to receive data with less than 1 ms queuing delays need to use a connection with a bandwidth of at least 800 Mbit/s (EMDI, All products) or 500 Mbit/s (EMDI, DAX® equities only or ETF only).

Enhanced Order Book Interface market data is currently only available to Trading Participants using 10 Gbit/s connections.

Trading Participant are advised to take two cross connects (one for each market data stream) in Co-Location to receive market data.

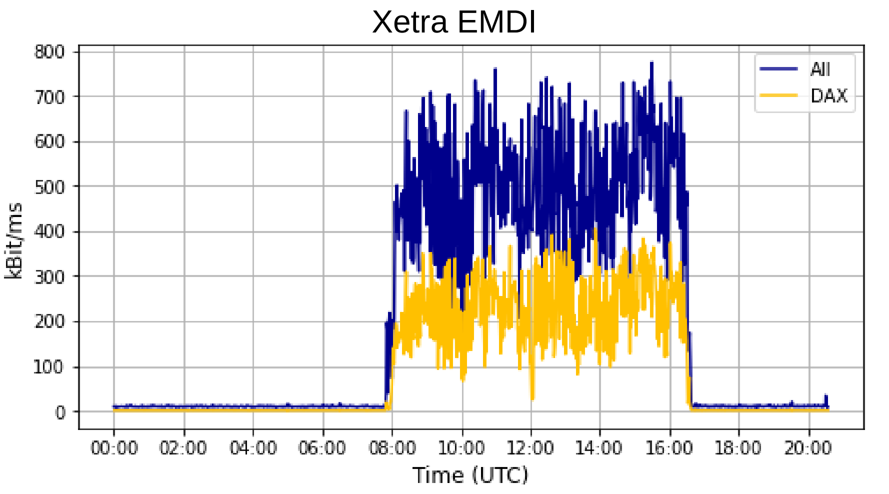

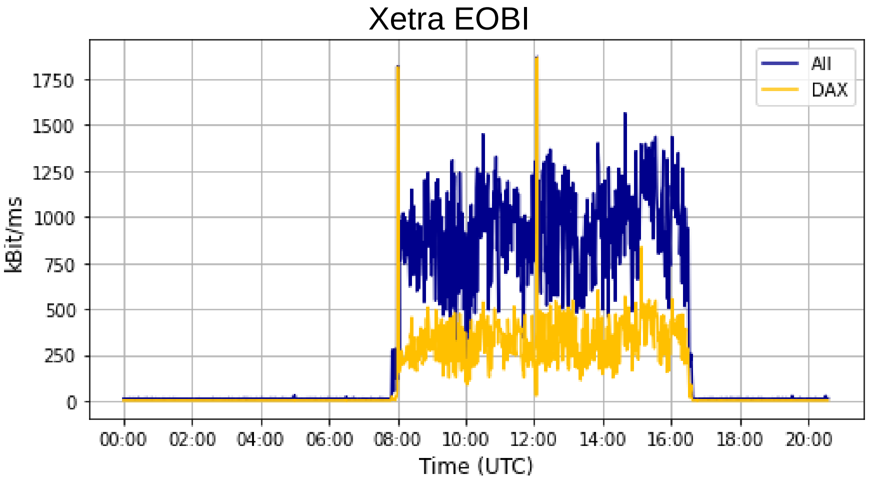

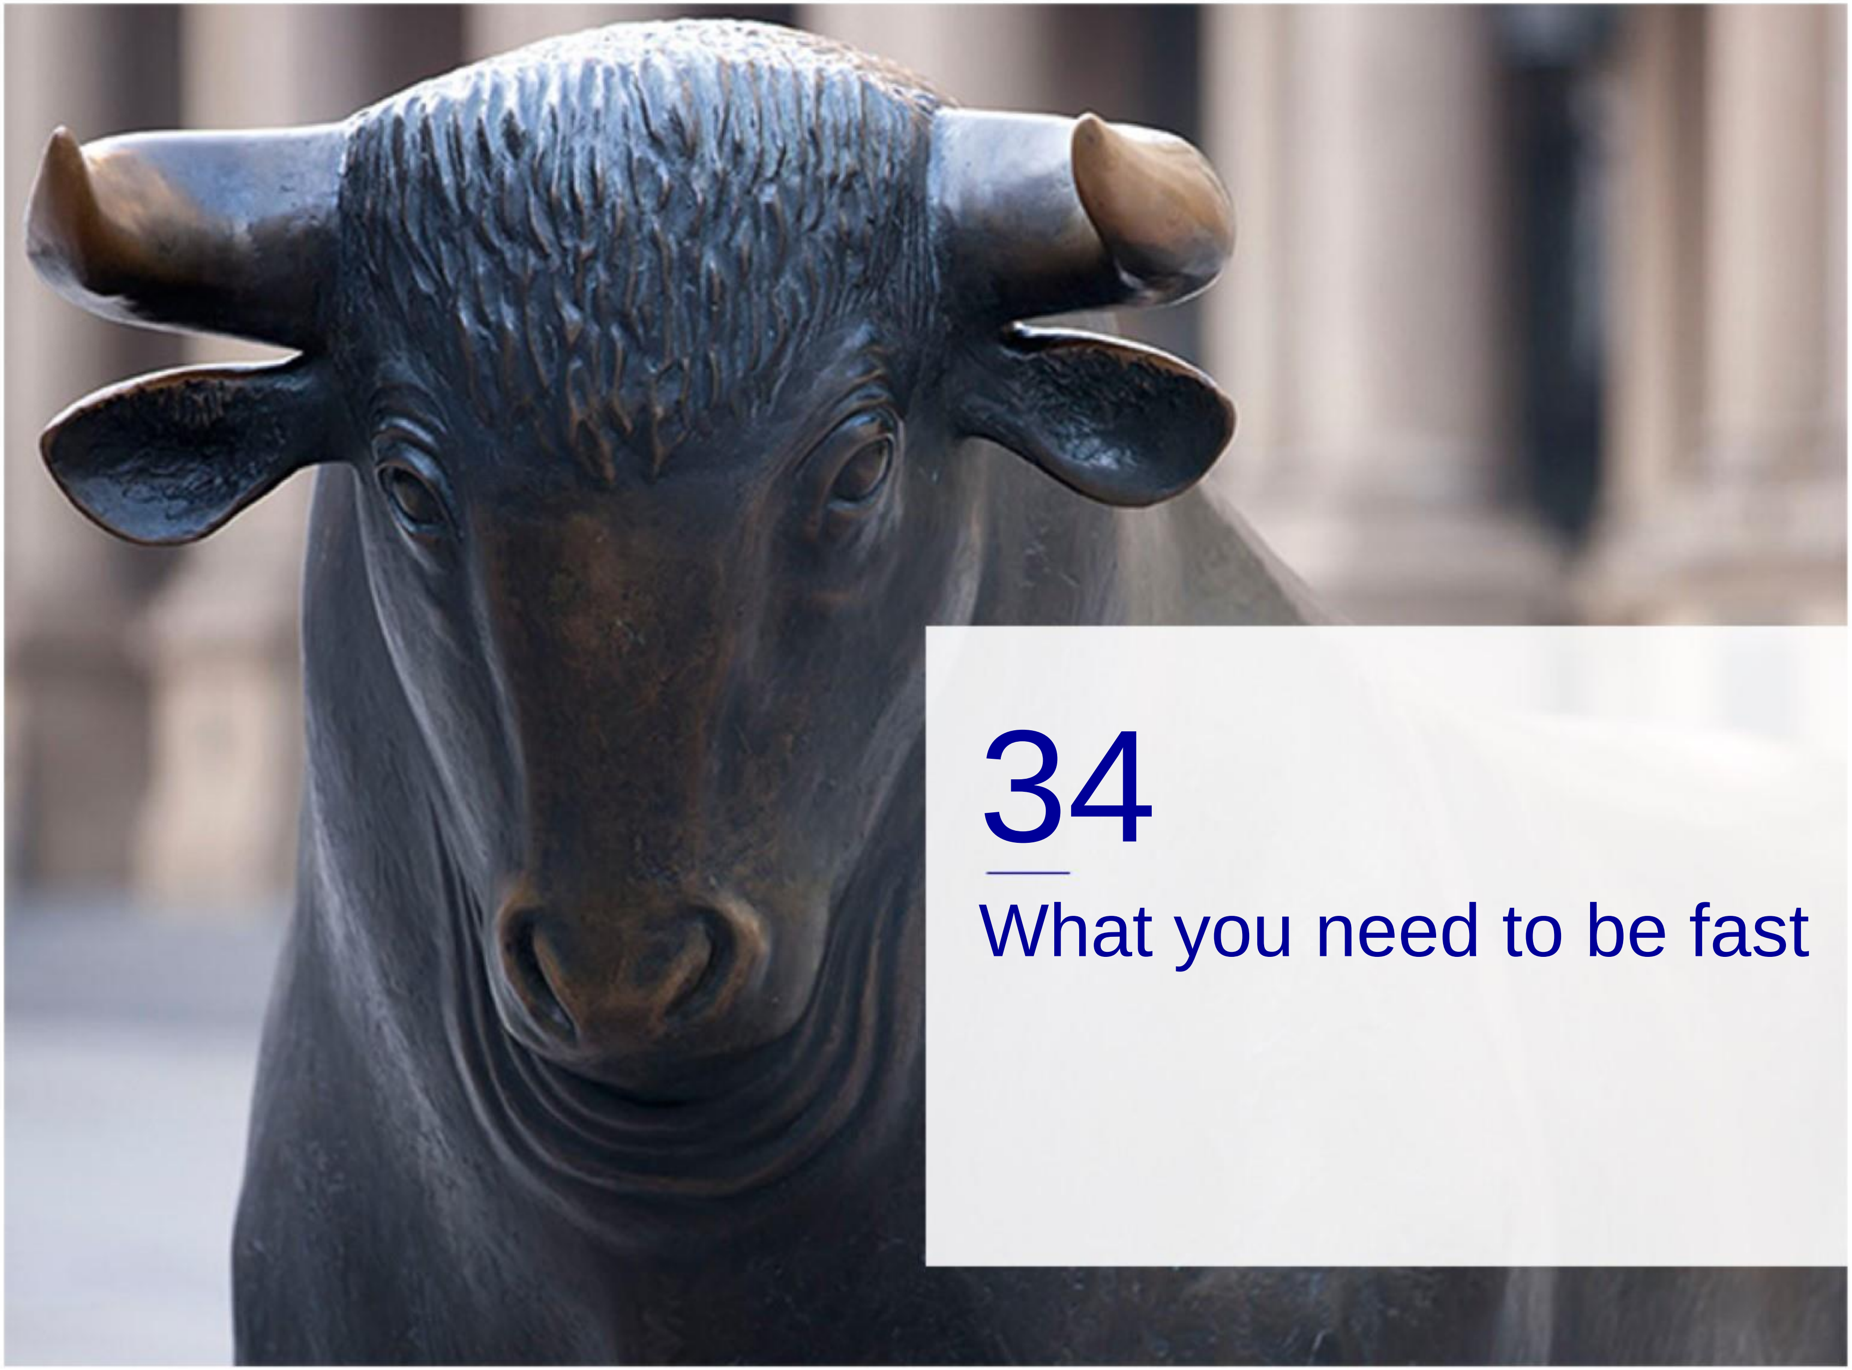

34

What you need to be fast

What you need to be fast…

A few recommendations to achieve the low latency

Use the Equinix Co-Location facility to be close to Deutsche Börse T7.

Use state-of-the-art switches (if any) and only have at most one hop between the exchange network and your server.

Use good network interface cards and TCP/IP acceleration, e.g. a kernel-by-pass library.

Use at least two 10 Gbit/s (cross-) connections in Co-Location for EOBI market data and two 10 Gbit/s connections for T7 ETI.

Use PS gateways and make sure you use the cross connect on the same side as the gateway you are connecting to (compare time-to-live values in the IP header in the responses from both sides).

Measure and analyze your own timestamps (e.g. the reaction time as recommended on the next slide).

Use state of the art time synchronization, i.e. GPS clocks and a high quality time distribution. The PTP signal you can get from us has a quality of +/- 50 ns. For our network timestamps we use the White Rabbit protocol and PPS breakouts. You can connect to our white rabbit time service providing you a time synchronization quality of 1-2 ns max, see https://www.deutsche-boerse.com/dbg-en/products-services/ps-technology/ps-connectivity-services/ps-connectivity-services-time-services

We provide highly accurate network timestamps of all orders leading to a market data update via the high precision timestamp file service, see https://datashop.deutsche-boerse.com/high-precision-timestamps.

Use the EOBI Execution Summary for fast trading decisions and position keeping (passive executions). For a consistent order book, all incremental updates following the Execution Summary should always be processed. For fastest decisions evaluate the market data classification based on the DSCP flags in the IPv4 header of EOBI market data packets.

Trade notifications need to be processed to create safety. We recommend to use either a low-frequency ETI session or a FIX trade capture drop copy to confirm the fast execution information provided by the execution reports via high-frequency sessions.

What you need to be fast…

Participant reaction time measurement

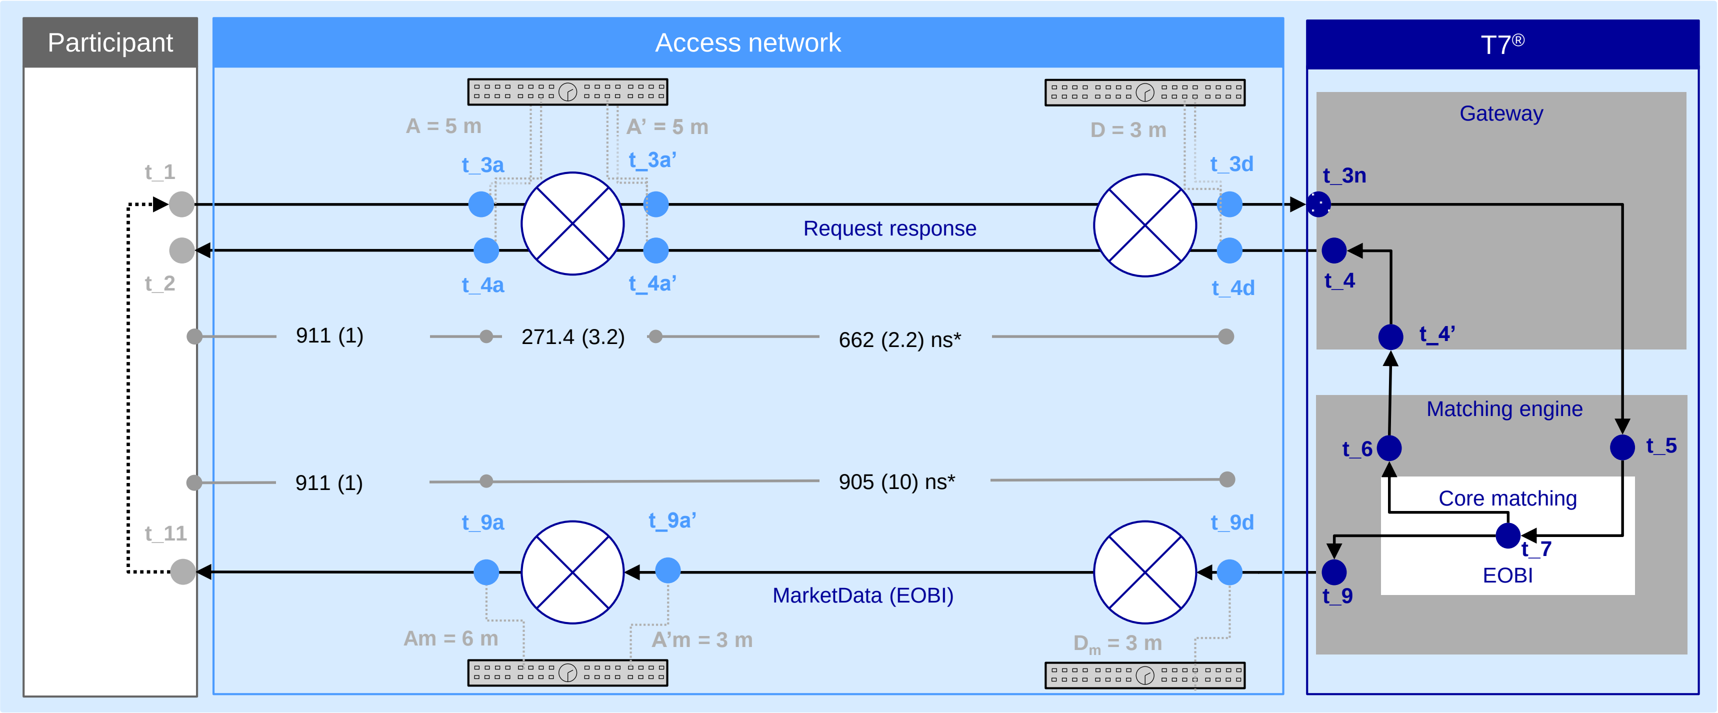

Measure the time between market data reception (t_10/t_11) and your reaction (t_1) and align with timestamps from the high precision timestamp file

*latencies are measured in nanoseconds given as 'median (variance)'
the represent tap to tap latencies NOT including the cable tap to timestamping device, thus recorded latencies may deviate because of different tap to measuring device cable length

t_3a, t_3d and t_9d are available via the high precision timestamp file service, see https://datashop.deutsche-boerse.com/high-precision-timestamps
Take our white rabbit signal to compare your timestamps with ours with ns accuracy
https://www.deutsche-boerse.com/dbg-en/products-services/ps-technology/ps-connectivity-services/ps-connectivity-services-time-services#tab-1556860-1556864

37

T7 Overview

Overview of T7

- T7 consists of partitions. A partition is a failure domain in charge of matching, persisting and producing market data for a subset of products. Each T7 partition is distributed over two rooms in the Equinix data centre.
- There are 10 Eurex T7 and 10 Xetra T7 partitions.
- Separate partitions are used for markets of other exchanges hosted on T7 (e.g. Vienna (XVIE), EEX (XEEE), Bulgaria (XBUL), …).
- The reference data contains the mapping of products to partition IDs.
- With the introduction of PS gateways there is a one-to-one mapping of active PS gateways to partitions. The default active PS gateways are located on the same side as the active matching engines.
- 4 LF gateways allow access to all Eurex partitions and the separate EEX partition.
- 4 LF gateways are shared between all cash markets (XETR, XVIE, XBUL,…).

- Note that the active half of a partition and its PS gateway is either on side A (for even partitions) or on side B (for odd partitions).
- Only in case of the failure of a matching engine or a market data publisher, the active half of the service will shift to the other room.
- With consolidated matcher/gateway processes (Q2 2021) the active PS gateway and active Matching Engine act as a single failure domain within each partition, i.e. they will always fail as a single logical group.

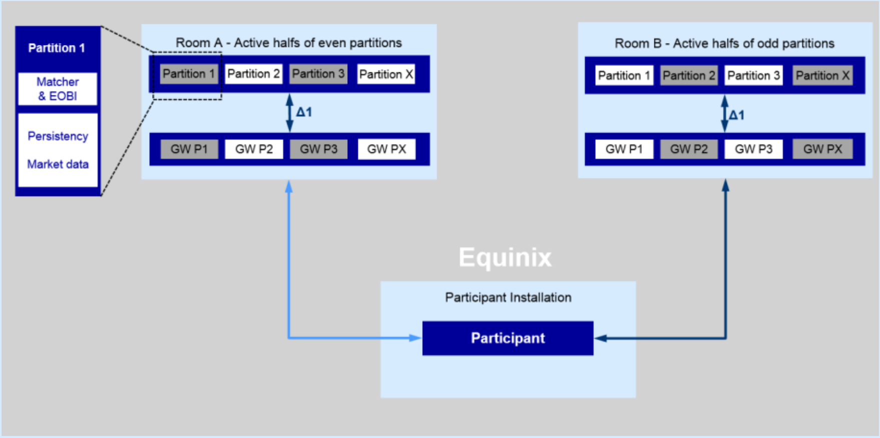

T7® Topology

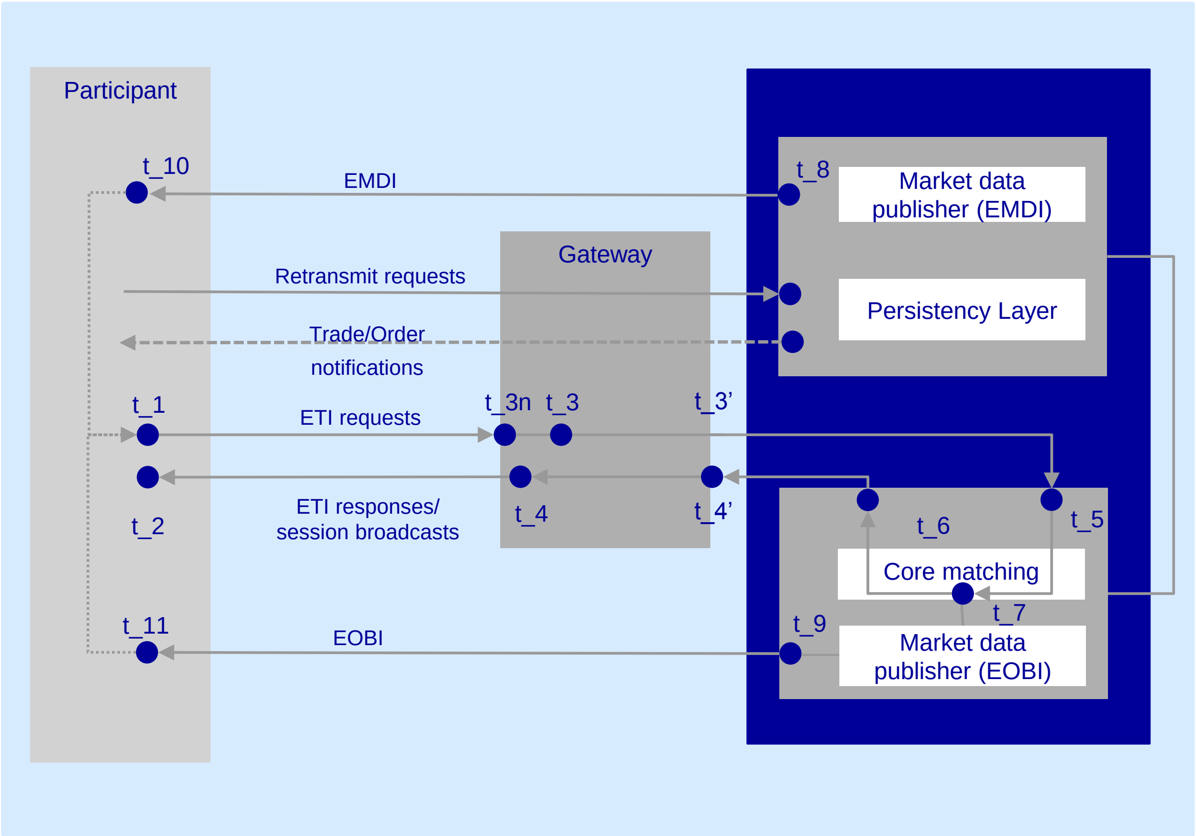

Matching engine:

- order book maintenance & execution
- creation of direct responses as well as execution messages for passive orders/quotes
- creation of EOBI order book messages
- creation of EOBI order book snapshot messages

Persistency:

- persistent order storage
- trade/execution history
- transaction history for standard orders
- creation of listener broadcast for standard orders

Market Data (EMDI):

- creation of order book delta messages
- creation of order book snapshot messages

Partitions
Overview

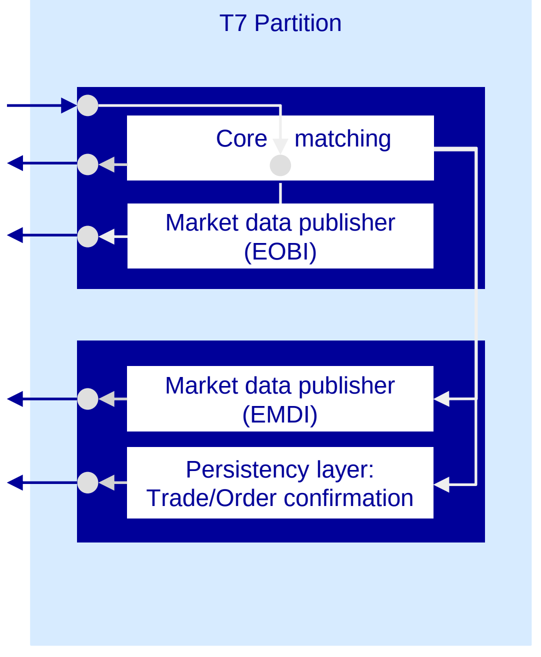

Orders/quotes entered for a specific product are sent by the gateway server to the respective matching engine (residing in a partition).

The matching priority is assigned when the orders/quotes are read into the matching engine.

The core matching component works as follows:

- when an order/quote arrives, it is functionally processed (e.g. put in the book or matched),
- handover of data to the EOBI market data publisher and
- handover of all data resulting from the (atomic) processing of the incoming order/quote to the market data and persistency components in the partition.

Resulting responses and private broadcasts are sent out in the following order:

- direct response to the order/quote entered (for persistent as well as for non-persistent orders and quotes),
- fast execution information for booked orders/quotes (in case of a match).

The generation of listener broadcasts, trade confirmations (by the persistency server) and of non-EOBI market data (by the market data publisher) is done on separate servers. Hence the order of the resulting messages from these servers is not strictly deterministic.

Note that the matching engine holds states of orders in memory. All responses, broadcasts and EOBI market data thus are preliminary in nature.

Order Entry
Overview

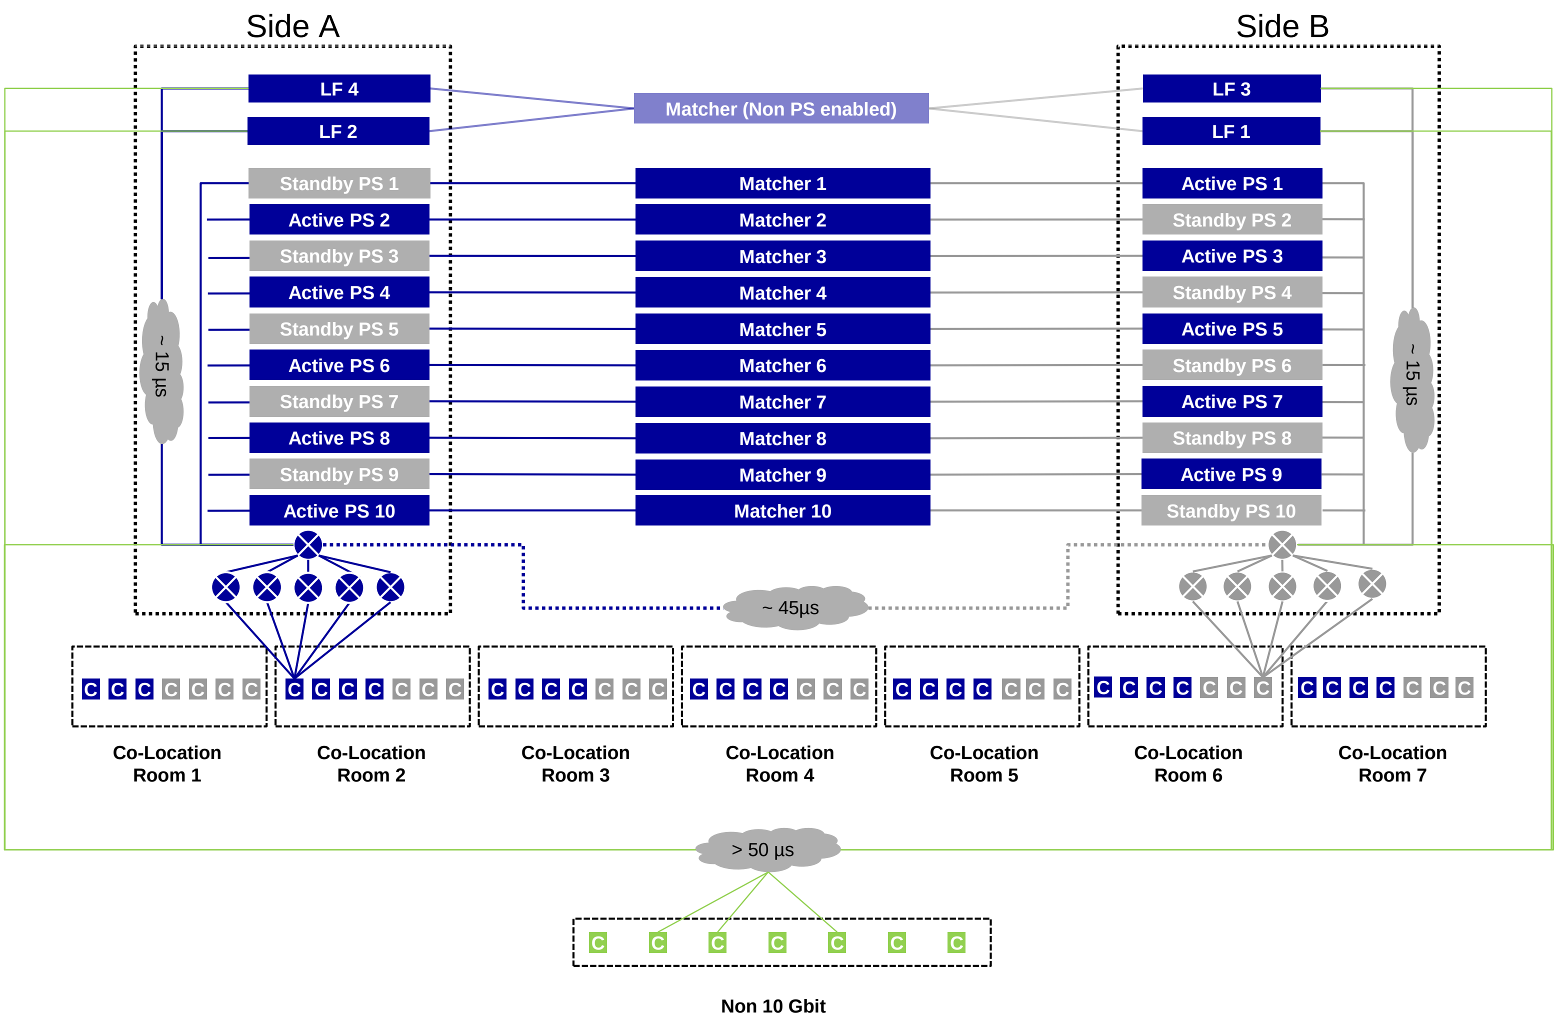

Order Entry
Introduction

Requests sent to T7 will be routed via an access network and a gateway.

There are the following basic connection alternatives:

Choice of Network

There are two network classes connecting a participant's installation with the T7 gateways:

- Co-Location with 10 Gbit bandwidth and a one way base latency of around 2 µs.
- Other networks with less than 1Gbit bandwidth and a one way base latency of around 50 µs.

Choice of Session Type

T7 supports three session types:

- High frequency sessions connect to PS gateways for low latency access to a single partition (binary ETI protocol).
- Low frequency sessions connect to LF Gateways for convenient access to all partitions, with a considerably higher base latency (ETI).
- FIX sessions connect to FIX gateways for convenient access to all partitions using the FIX protocol, with a considerably higher base latency than LF gateways.

Important Developments

- LF gateways moved closer to 10 Gbit network and PS gateways (December 2019 / January 2020).
- LF gateway requests are routed via PS gateways (February 2020), see page 7.
- LF gateway tuning measures lead to a decrease of processing time (March 2020).

Order Entry
Co-Location Network

Participants may use Co-Location to place their infrastructure in the datacentre that hosts the T7 system.

The Co-Location 10 Gbit network has the following properties:

Fair and equal access

Regardless of the Co-Location room we ensure all lines are created equal.

More precisely the latency between the handover point in the participant's rack and the first (access layer) switch is calibrated to below +/- 2.5 ns.

Two redundant halves ('A' and 'B')

There are two independent order entry network halves.

As active gateways are placed in either half there is an optimal side for each gateway (even numbered gateways are on the A side, odd number on the B side).

Crossing sides, e.g. connecting to a B side gateway via an A network, is possible but results in a 45 µs higher base latency.

Two hierarchical switch layers

Participants connect to access layer switches (currently 5 Eurex*, 2 Xetra per side).

The uplink of each access layer switch is connected to a distribution layer switch.

The distribution layer switches have a direct connection to the active gateways on the respective side.

* For capacity reasons a 6th access layer switch will be added on Eurex B Side on 29 March 2021.

Order Entry
Gateways

There are three gateway types to access the T7 system:

Partition-specific (PS) Gateway

Protocol:	flat binary (ETI)
Allowed session types:	High Frequency Sessions only
Sequencing:	FIFO operation (Sequence guaranteed from network card to matching engine in)
Latency:	lowest, median latency ~ 11 µs network card to matching engine in
Versatility:	Allows routing to one partition only, only subset of broadcasts available

Low Frequency (LF) Gateway

Protocol:	flat binary (ETI)
Allowed session types:	Low Frequency Sessions only
Sequencing:	FIFO not guaranteed
Latency:	medium, (additional 43 µs median latency compared to PS gateway direct access)
Versatility:	Routes to all partitions (via PS gateway), all ETI broadcast types available

FIX Gateway

Protocol:	FIX
Allowed session types:	Fix Sessions only
Sequencing:	FIFO not guaranteed
Latency:	high, requests to the matching engine are routed via LF gateways
Versatility:	Allows routing to all partitions (via LF gateway), all FIX broadcast types available

Market Data
Overview

Market Data can be consumed over two distinct types of networks and in various types

Choice of Network

There are two network classes available for market data:

- Co-Location with 10 Gbit bandwidth and a one way base latency of around 2 µs.
  10 Gbit connections are equalized in length (cable latency difference of less than +/- 2.5 ns) and provide the lowest jitter.
- Other networks with less than 1Gbit bandwidth with higher base latency.

Choice of Market Data Type

There are three market data types:

- Order by Order market data (EOBI) with highest granularity and lowest latency in flat binary format.
  EOBI is sent out directly from the matching engine and is only available via 10 Gbit network.
- Price level aggregated market data (EMDI) with slightly higher latency in FAST encoded format.
- Netted price level aggregated market data (MDI) in FAST encoded format.

46

Appendix

Middleware, Network, Hardware and OS Overview

T7 uses state-of-the-art infrastructure components

Intel Xeon E5-2667 v3 CPUs (Haswell) on all servers hosting core services (Matching Engines, un-netted market data publishers)

Intel Xeon Gold 6144 CPU (Skylake) for PS gateways.

Intel Xeon Gold 6148 CPU (Skylake) on non-performance critical servers.

Servers hosting the consolidated Gateway/Matching Engine will be Xeon Gold 6256 (Cascade Lake).

We use the Red Hat Enterprise Linux operating system version 7.6.

T7 internal communication between its core components is based on Confinity Low Latency Messaging using an Infiniband network in order to deliver the required speed, capacity and stability requirements.

T7 network access
Deutsche Börse offers Trading Participants to connect via 10 Gbit/s cross connects to its T7 platform in the Equinix data centre.

The Co-Location offering uses Cisco 3548x switches operating in cut-through mode.
All cables are normalized to give an overall maximum deviation between any two cross connects of less than +/- 0.5 m(+/- 2.5 ns).
Insight into network dynamics is offered by the High Precision Timestamp File service (see https://datashop.deutsche-boerse.com/high-precision-timestamps).

Participant facing interface cards on the gateways and market data publishers use Solarflare EnterpriseOnload wire order delivery API to bypass the kernel TCP stack and deliver messages in the same order received by the network card.

Cables connecting Line-of-Sight antenna cables have been equalized to +/- 1m by Equinix.

Throttle and Session Limits

In order to protect the trading system, T7 has several measures in place to ensure that its most vital components are not harmed by a malfunctioning client application. Therefore transaction limits are imposed on T7 sessions.

All ETI sessions (HF and LF) are available with throttle values of 150 messages/sec or 50 messages/sec.
Furthermore LF sessions that cannot enter orders/quotes but can only receive trade and listener broadcasts are available (at a reduced price).

All ETI session types have an assigned disconnect limit of

- 450 for sessions with a throttle value of 150 messages/sec, i.e. a session will be disconnected in case of more than 450 consecutive rejects due to exceeding the transaction limit (throttle).
- 150 for sessions with a throttle value of 50 messages/sec, i.e. a session will be disconnected in case of more than 150 consecutive rejects due to exceeding the transaction limit (throttle).

Please note that in case the disaster recover facility is used, all ETI sessions will have a throttle limit of 30 messages per second.
For both limits, all technical transactions are counted using a sliding window.

The number of ETI sessions which can be ordered is limited. Currently, up to 80 sessions can be ordered. If more than 80 sessions are required please get in touch with your Technical Key Account Manager.

There is also a limit on the maximum number of sessions that can connect to a PS gateway concurrently per Participant. This limit is currently configured to 80 sessions, see Eurex Circular 122/17.

On 1 July 2019, we introduced a limit on the maximum number of outstanding session and trader login requests possible per business unit and per session at any given point in time. This limit is set to 50 on business unit level, 10 on session level. We recommend a synchronous login procedure, where a login request is sent on a session only after the previous login has been responded to. Please refer to the Incident Handling Guide for details.

For Eurex the number of order entry cross connects in colocation that may be used concurrently on a single day is limited to 36.
In addition the number of allowed ethernet frames per cross connect is currently limited to 25.000 per second and 300.000 per minute, see Eurex Circular 034/20.

T7 Topology
Timestamps

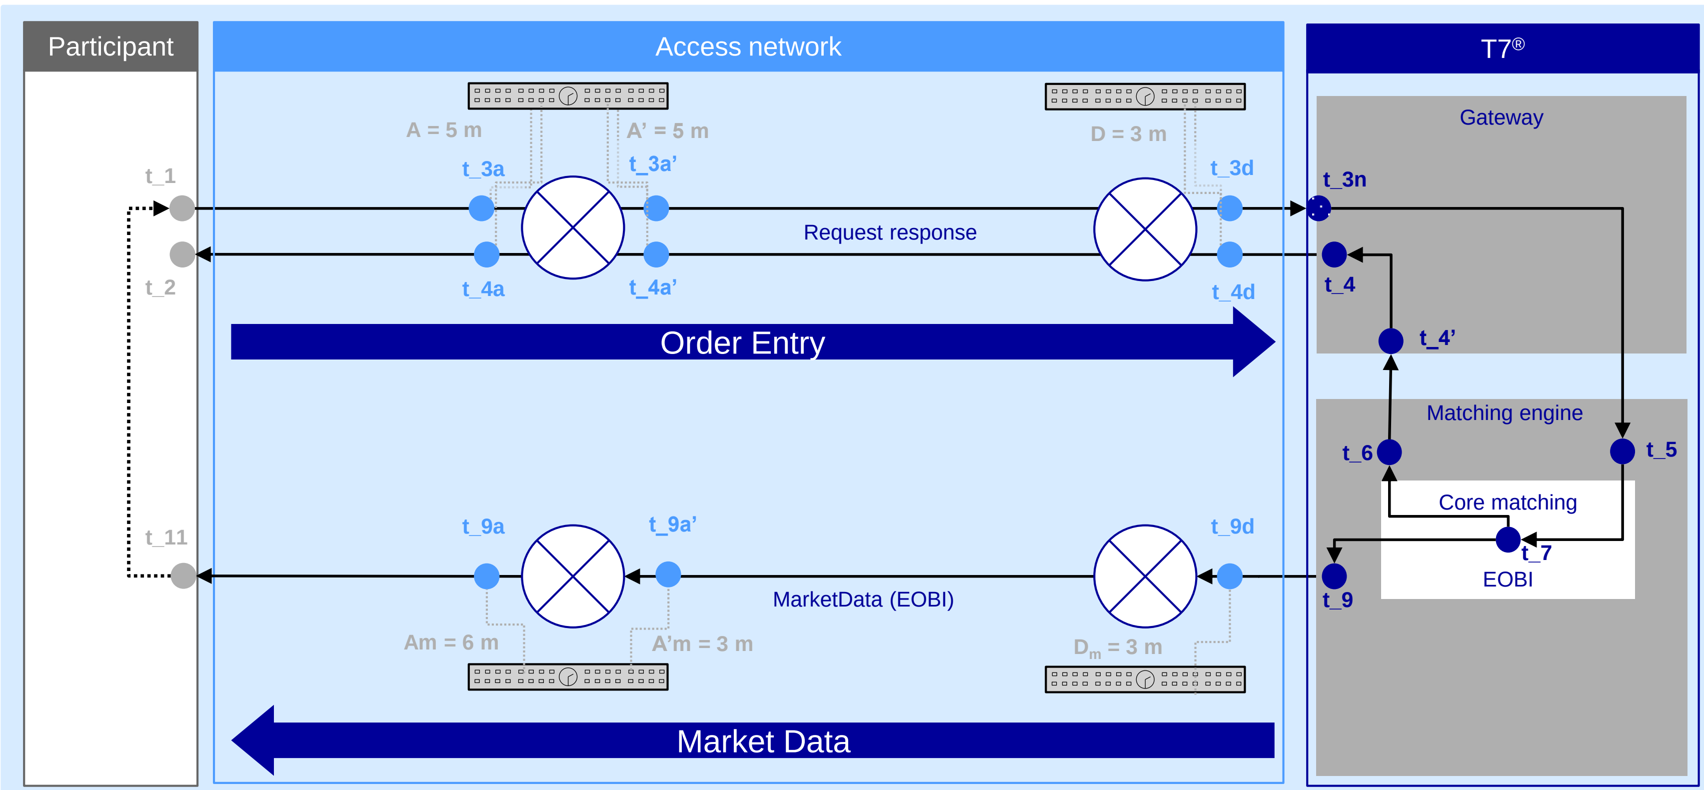

t_3a, t_3d and t_9d are available via the high precision timestamp file service, see https://datashop.deutsche-boerse.com/high-precision-timestamps
Take our white rabbit signal to compare your timestamps with ours with ns accuracy https://www.deutsche-boerse.com/dbg-en/products-services/ps-technology/ps-connectivity-services/ps-connectivity-services-time-services#tab-1556860-1556864

T7 timestamp reference

The timestamps t_3,…,t_9 are available via the following fields:

t_3, t_3n:	Tag	5979	(“RequestTime”)	in the T7 ETI Response in the T7 EMDI Depth Incremental message, in case a trade is reported in the T7 EOBI Execution Summary message
	Tag	21002	(“TrdRegTSTimeIn”)	in the T7 EMDI Depth Incremental message, in case a trade is reported in the T7 EOBI Execution Summary message
t_4’:	Tag	7765	(“ResponseIn”)	in the T7 ETI Response (from the matching engine)
	Tag	25043	(“NotificationIn”)	in the T7 ETI Notification (from the matching engine)
t_4:	Tag	52	(“SendingTime”)	in the T7 ETI Response and Notification
t_5:	Tag	2445	(“AggressorTime“)	in the T7 EMDI Depth Incremental message, in case a trade is reported in the T7 EOBI Execution Summary message
t_6:	Tag	21003	(“TrdRegTSTimeOut“)	in the T7 ETI Response and Notification (from the matching engine)
t_7:	Tag	17	(“ExecID“)	in the T7 ETI Response (from the matching engine) in the T7 EOBI Execution Summary message
	Tag	273	(“MDEntryTime”)	in the T7 EMDI Depth Incremental and Top of Book Implied message
	Tag	60	(“TransactTime”)	in the T7 EMDI messages for other events
	Tag	21008	(“TrdRegTSTimePriority”)	in the T7 EOBI Order Add and Order Modify messages
	Tag	60	(“TransactTime”)	in the T7 EOBI Order Modify Same Priority and Order Delete messages
t_8:	no Tag		(“SendingTime”)	in the T7 EMDI UDP packet header
t_9:	Tag	60	(“TransactTime”)	in the T7 EOBI packet header
(t_8 - t_5):	no Tag		(“PerformanceIndicator“)	in the T7 EMDI UDP packet header of the T7 EMDI Depth Incremental stream

Notes on timestamps:
All timestamps provided are 8 byte integers (in nanoseconds after Unix epoch).
The PerformanceIndicator is a 4 byte integer (in nanoseconds).
The Network timestamps (t_[x]a, t_[x]a’, and t_[x]d, t_[x]d’) are not available in any protocol field but some via the High Precision Timestamp File service.

Thank you for your attention

Contact

Sergej Teverovski, Manfred Sand & Phuong Hieke

E-Mail monitoring@deutsche-boerse.com

For updates refer to https://www.eurex.com/ex-en/technology/t7

http://www.xetra.com/xetra-en/technology/t7/publications