Eurex Exchange’s trading system

Some insights into the details that matter for high frequency trading!

April 2012
W. Eholzer

eurex

Agenda

- Achievements
- Topology
- Inside the matching engine
- Infrastructure
- Outlook
- Appendix

Eurex Technology Roadmap

In 2006 Eurex initiated its Technology Roadmap to deliver innovative and superior technology solutions.

This initiative includes:
- provision of different interfaces that meet the needs of different user groups,
- high throughput and low latency of the trading system,
- exceptional level of transparency & customer service.

Eurex Group is very committed to these goals and is constantly improving its technical offering.

The Eurex Technology Roadmap will continue in 2012 with the next phase, whereas from Q4 the report distribution via the MISS will be discontinued.

Furthermore, Eurex Exchange will introduce a new trading architecture starting December 3, 2012.*

The initial roll-out is followed by a migration phase where products will be moved in a stepwise approach to the new trading architecture. The new system is based on Deutsche Börse Group's proprietary global trading architecture, which is already in use at the International Securities Exchange (ISE).

For further details please visit Eurex circulars 183/2011 and 21/2012 and have a look at our web page about the new trading architecture: www.eurexchange.com/nta

** Note that the launch of the new trading architecture is subject to the required legal approvals.*

Futures order add round-trip times

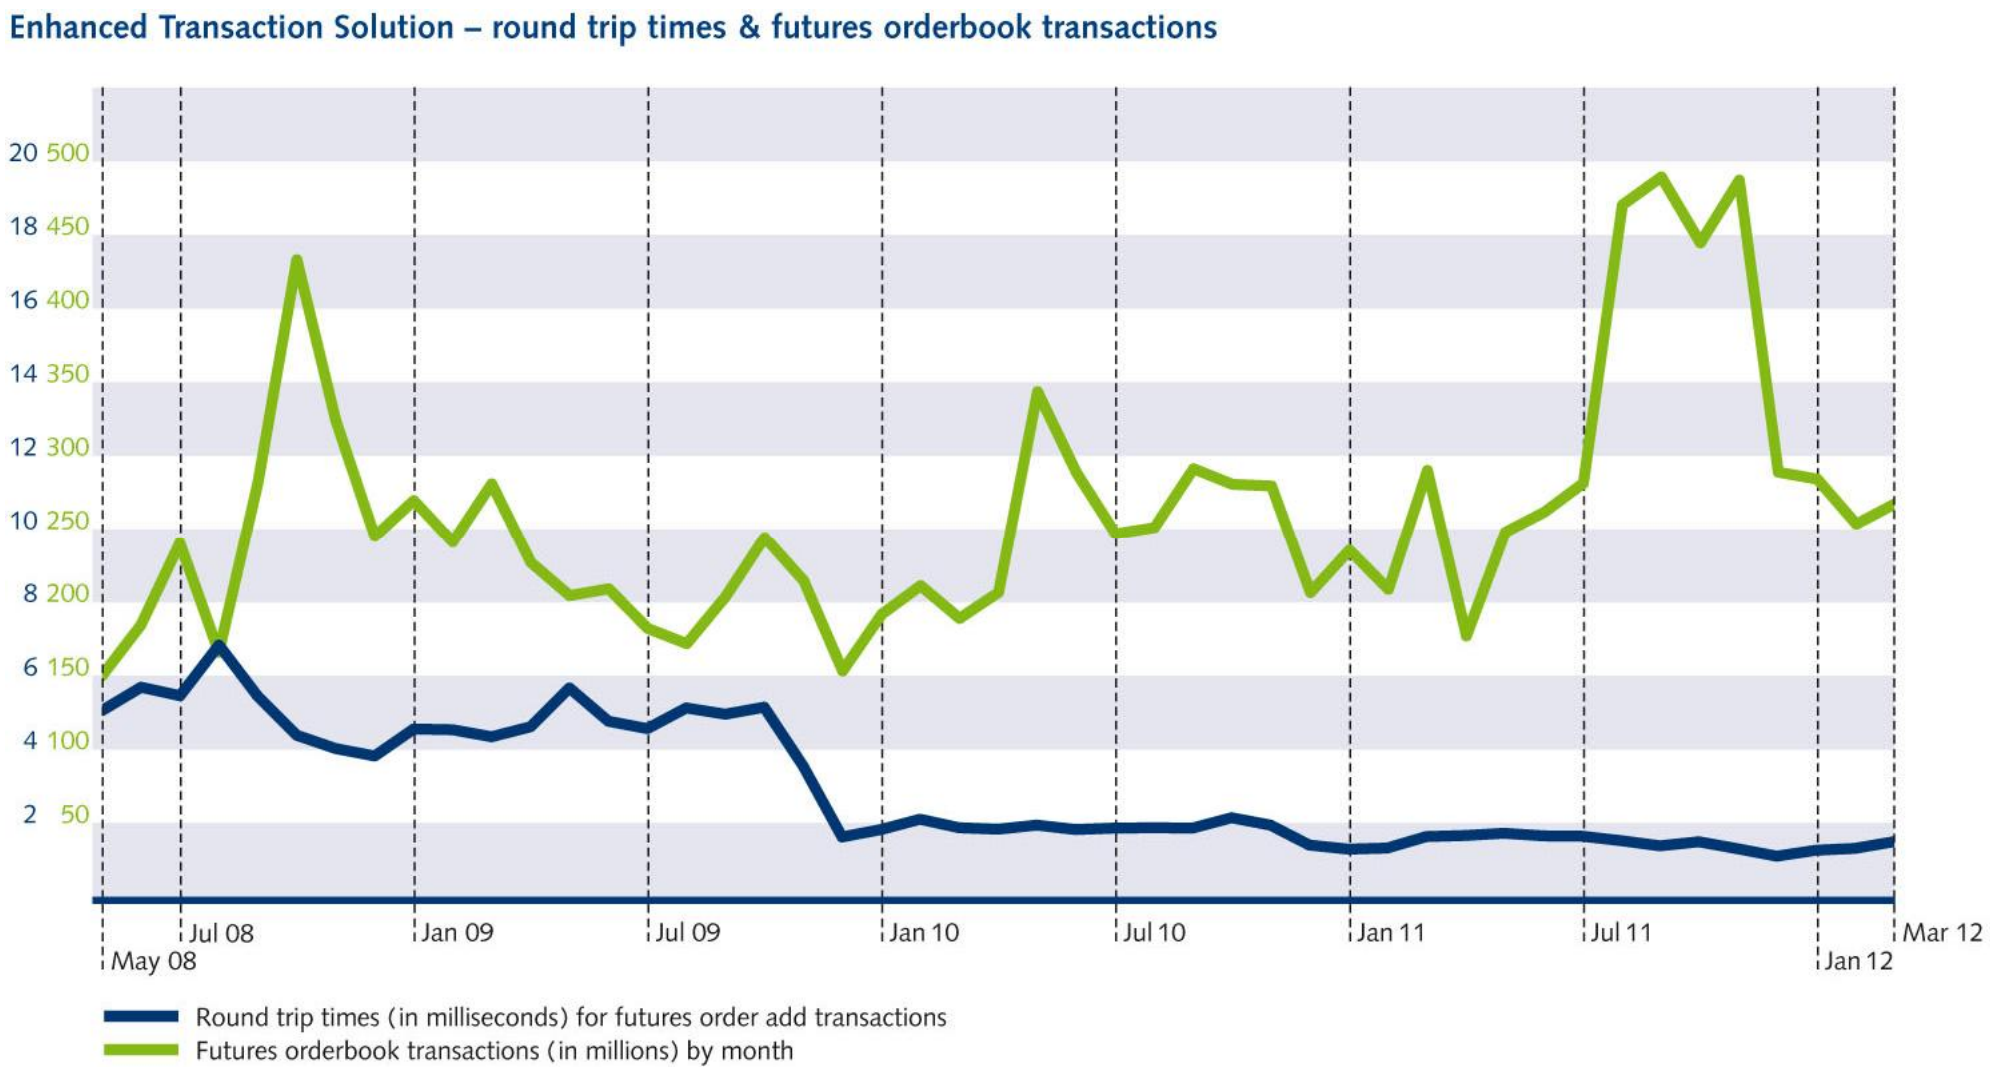

Futures order add latency has been significantly reduced to a level of a daily average of about 0,8 -1,2 ms for users of the Enhanced Transaction Solution located in the Equinix data center in Frankfurt.

Usage of the Enhanced Transaction Solution

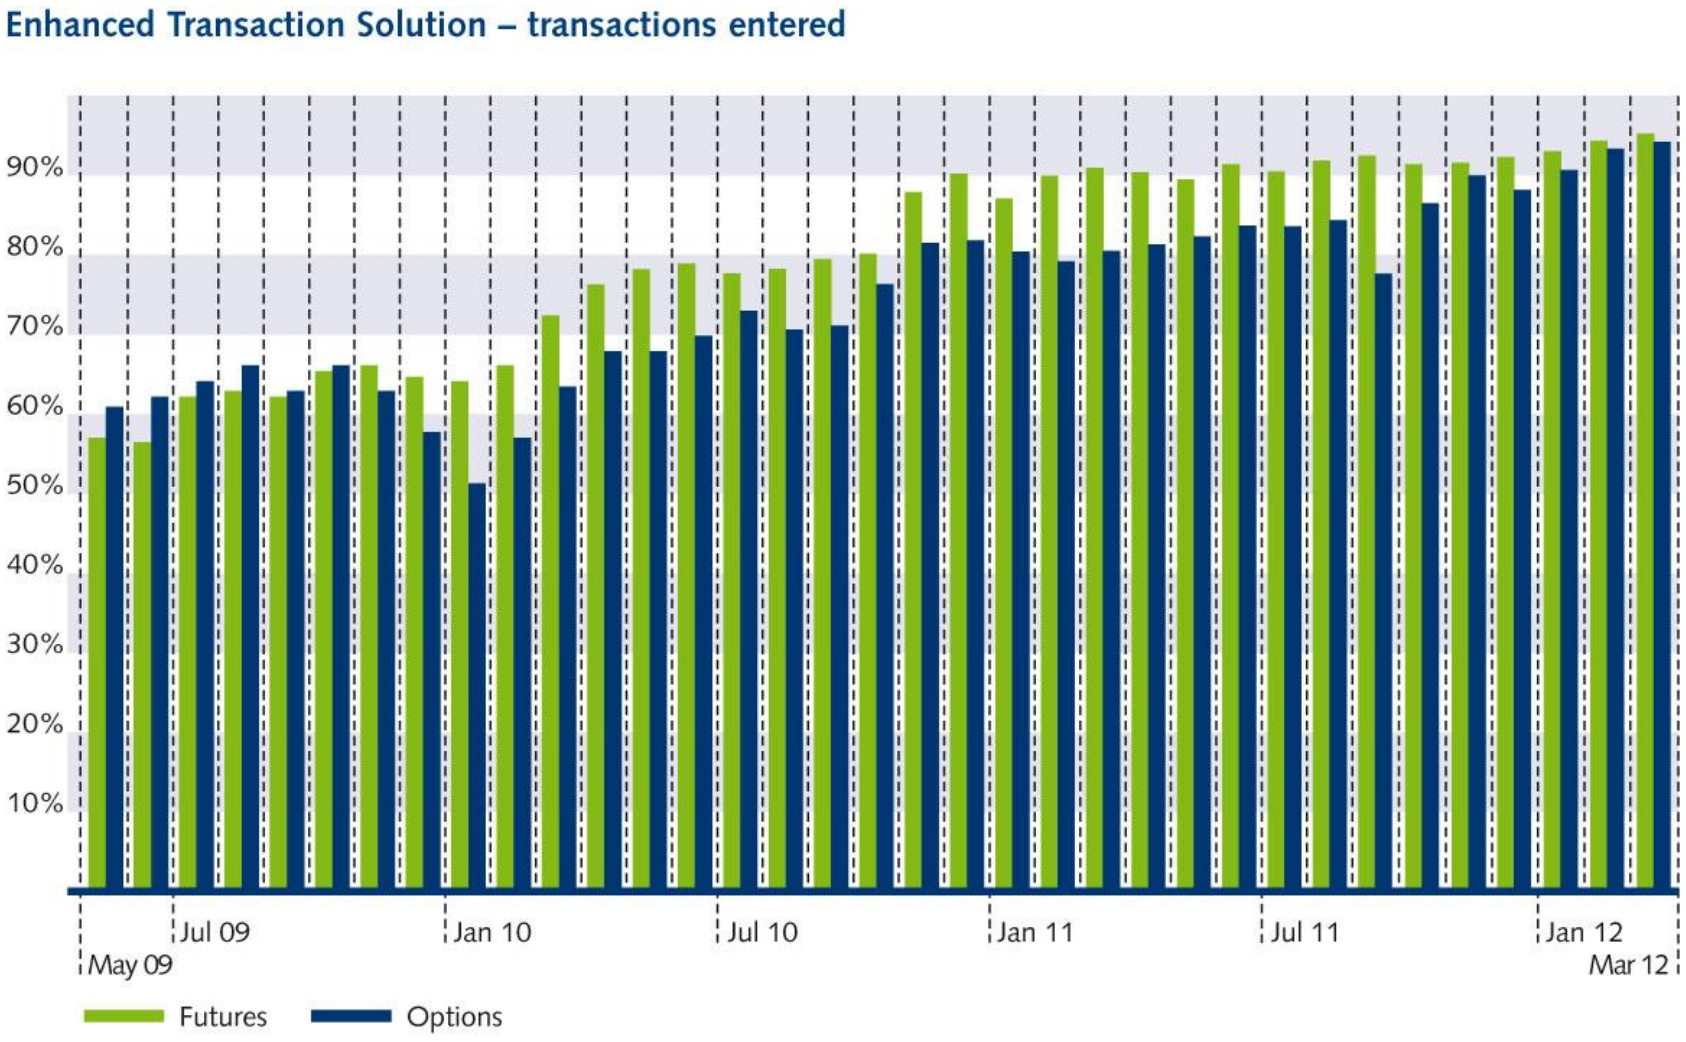

Today about 95 per cent of options and futures order and quote transactions are entered using the Enhanced Transaction Solution.

Processed transactions & response times

Number of processed transactions at Eurex Exchange & response times

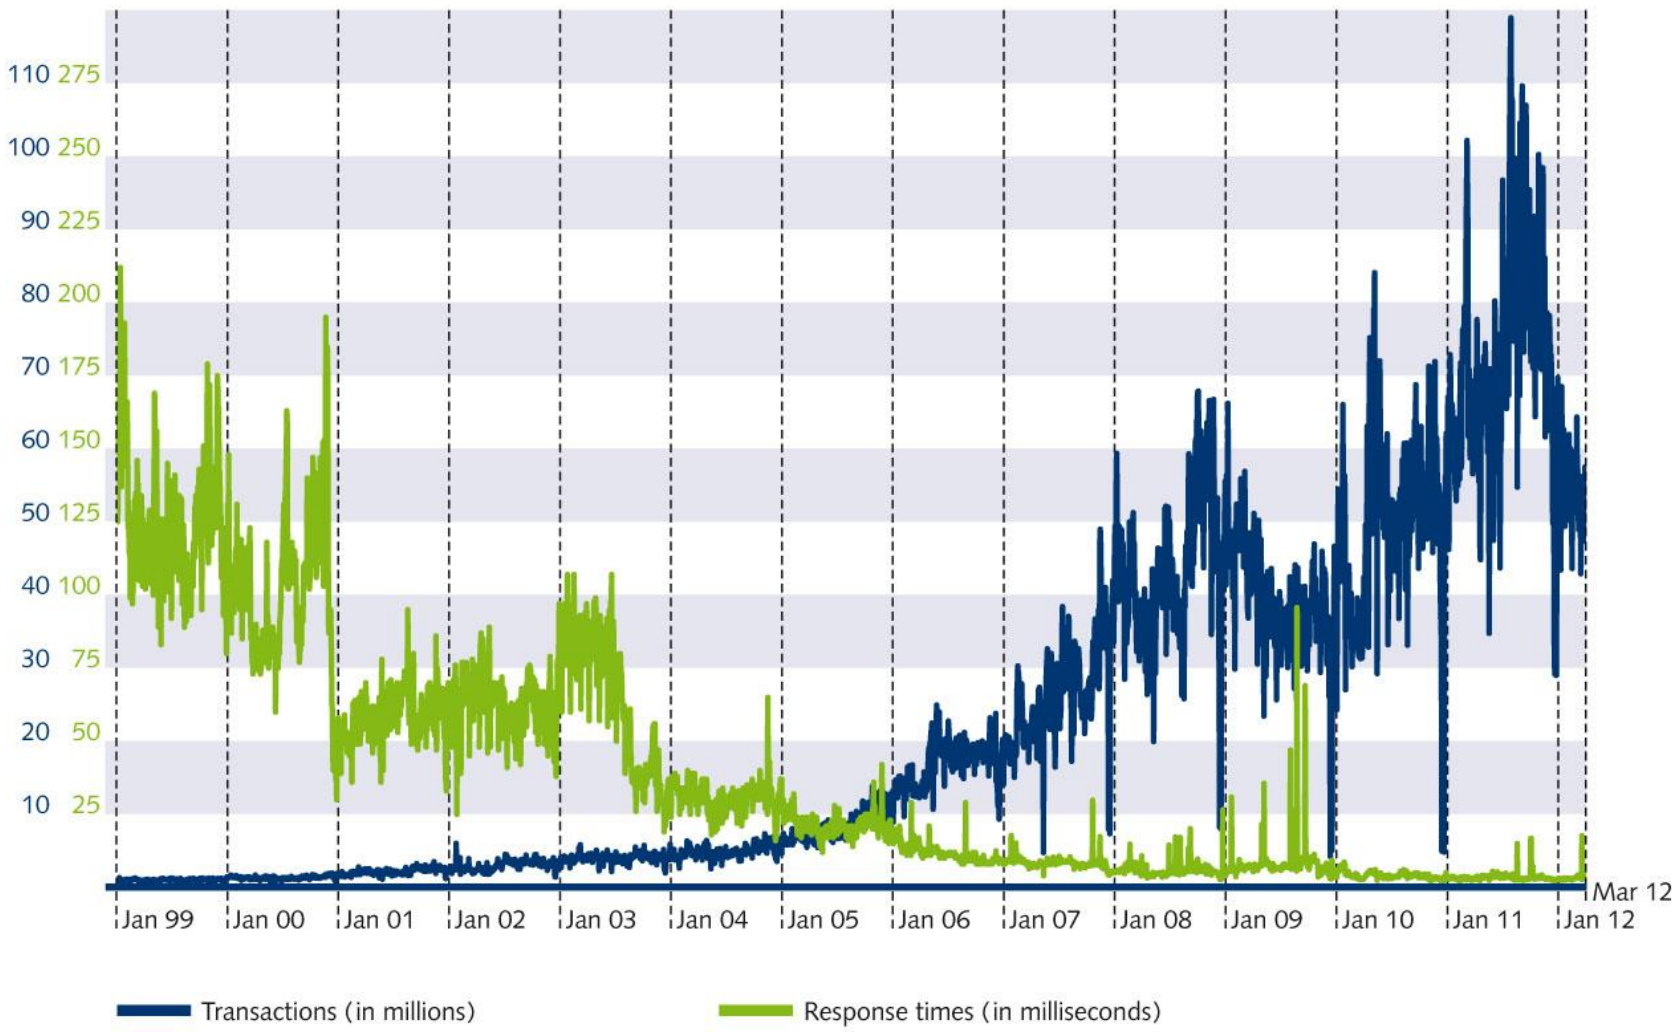

Eurex has continuously invested in its trading system and has been able to reduce the processing time of technical transactions significantly although the daily load on the system has grown extremely.

Agenda

- Achievements
- Topology
- Inside the matching engine
- Infrastructure
- Outlook
- Appendix

Overview

- The 14 Eurex Exchange matching engines* are distributed over two rooms in the Equinix data center.
- The Enhanced Transaction Solution gateways in the Equinix facility provide the fastest access to the Eurex trading system. There are 12 such gateways in the Equinix data center shared by all Eurex trading participants.
- Eurex's Enhanced Transaction Solution reference data contains the OptiGatewayLocId, providing an exceptional level of transparency & customer service (the OptiGatewayLocId is mainly relevant for market data; details see below).

Data center topology

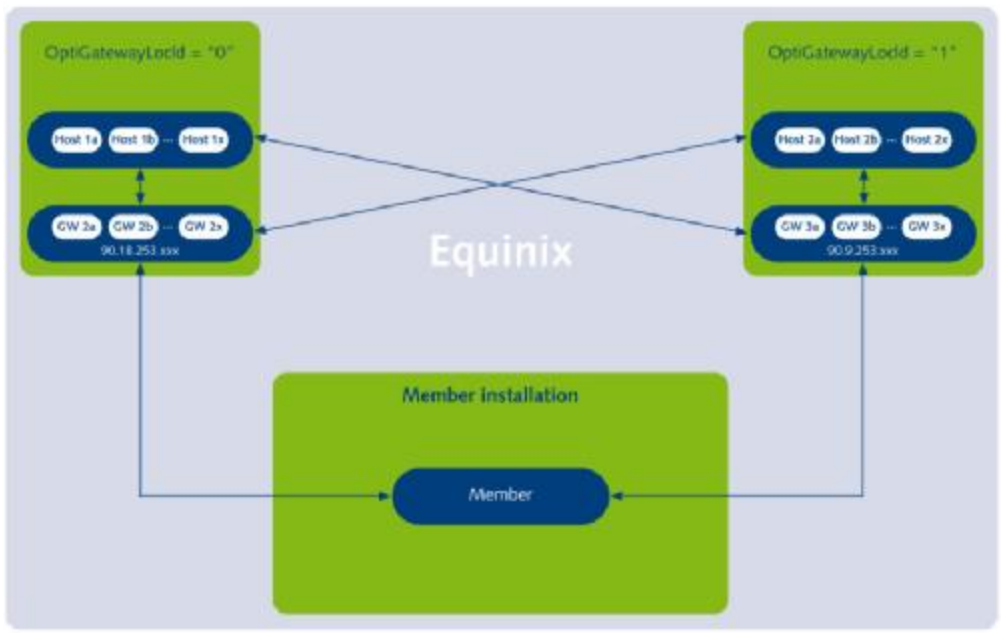

Eurex Exchange has increased the number of matching engines from 12 to 14 (from March 6,2012 onwards).

Currently we have the following mapping for our benchmark futures:

OptiGatewayLocId = "0":
FGBM, FGBS, FDAX, FSMI

OptiGatewayLocId = "1":
FGBL, FGBX, FESX

Note that the product to data center mapping (Equinix) typically changes on a time scale of weeks/month. In case of e.g. a hardware failure, the mapping can change intra-day (intra-day changes are not published).

Enhanced Broadcast Solution: topology

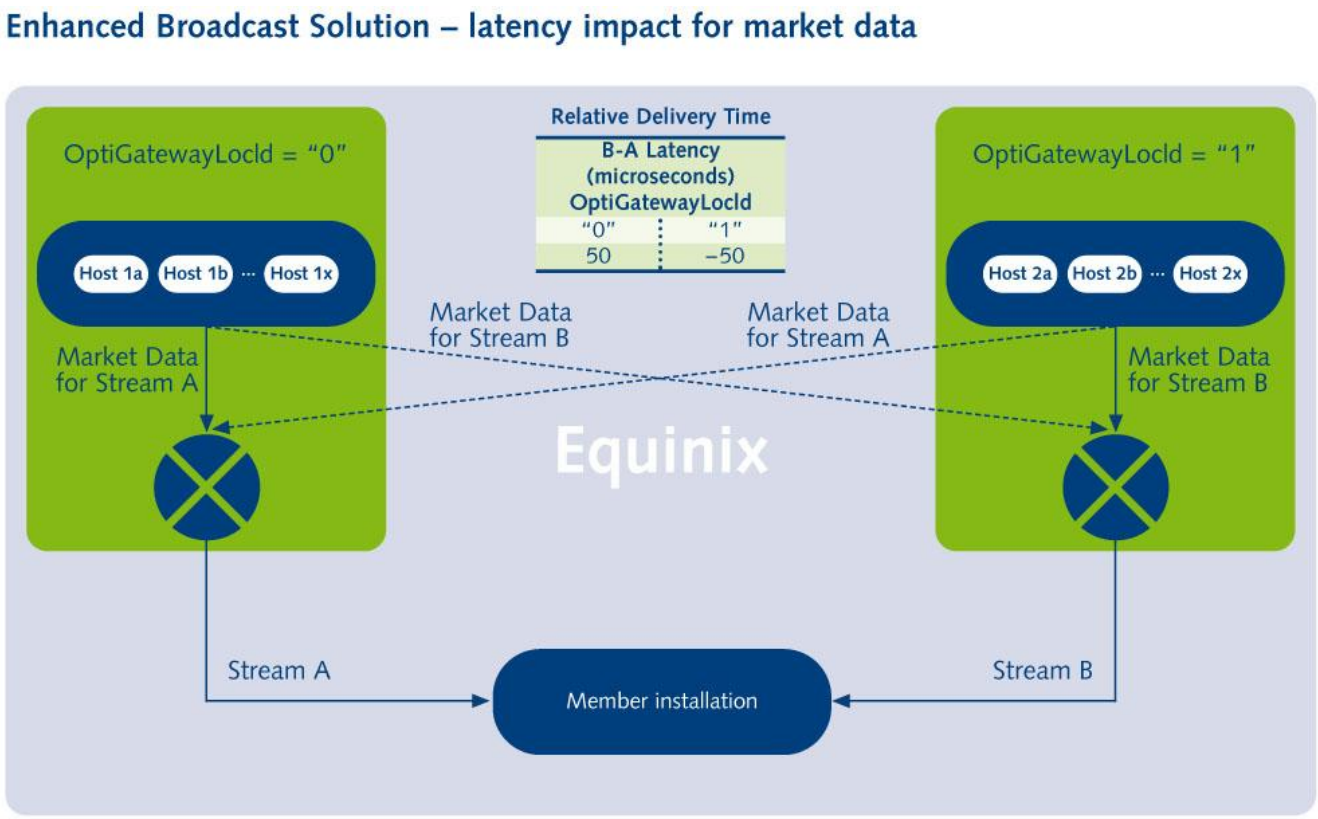

- Please note that for products with OptiGatewayLocID = "0" market data is published first on the A stream and then on the B stream whereas for products with OptiGatewayLocID = "1" market data is published first on the B stream and then on the A stream.
- The SourceID provided with each Enhanced Broadcast Solution message tells you on which matching engine the message was generated.

- Eurex provides a csv download file on a daily basis with the minute by minute network latency (minimum, average, maximum 99 per cent, 99.5 per cent) for the A and B stream of the Enhanced Broadcast Solution. This information can help you in case of operational issues with market data latency, if you had an issue or Eurex.
- Please note that the latency advantage of the 10 GE connections compared to 1 GE connections is approximately 60 usec for order book incremental messages.

Enhanced Broadcast Solution: data volume

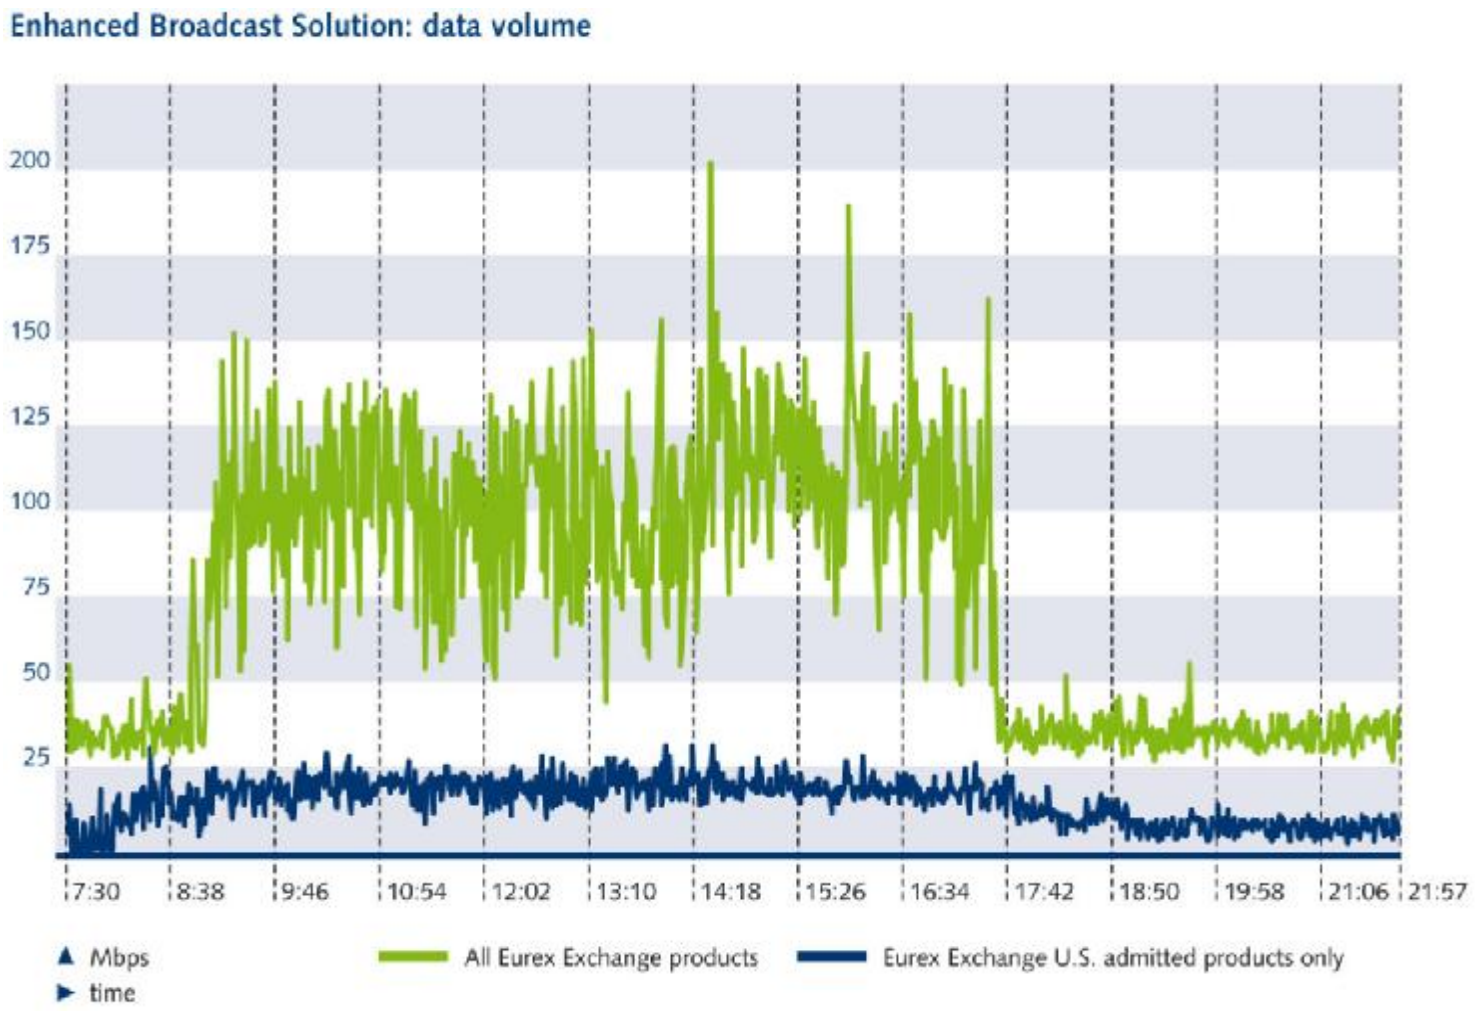

- The provided data shows one data point per minute for February 3, 2012.
- Each data point equals the maximum bandwidth produced on a 10ms scale by the complete A stream in kbps.
- Eurex Enhanced Broadcast Solution peak data volume has been about 20 per cent higher than on the sample day above. Hence participants that want to receive data for all Eurex Exchange's products or U.S. only products with less than 10ms queuing delays need to use a connection with a bandwidth of more than 250 Mbps (all products) or 40Mbps (for U.S. only products), respectively.

Enhanced Transaction Solution – optimal access

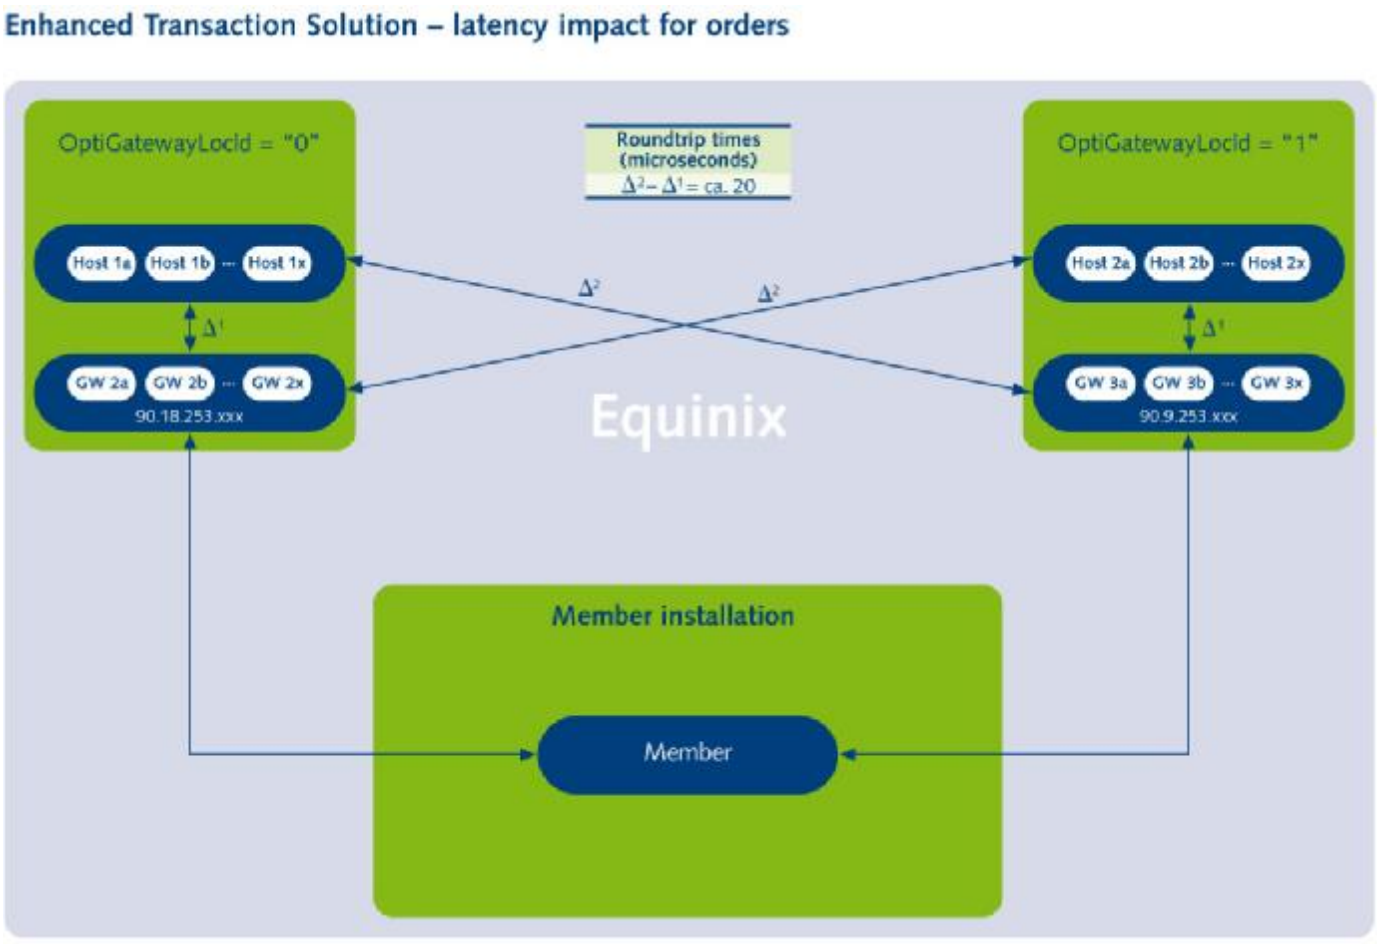

- Fastest access to Eurex Exchange is provided via the Enhanced Transaction Solution gateways in Equinix (see IP subnets in the diagram).
- Note that for optimal routing of orders the OptiGatewayLocID is only of marginal importance as the network related time difference of about 20 usec is much smaller than the variations of the processing times on the gateways (see slide 13 for details).

- Daily statistics about private "last mile" performance between the gateways and participant servers as well as best in class numbers (per location and system wide) are provided. A good daily average TCP/IP round trip time is 50 usec for 10 Mbit/sec connections and about 10 usec for 10 GE connections.
- Please note that the overall latency advantage of the 10 GE connections compared to 1 GE connections is approximately 50 usec round trip times for order transactions.

Enhanced Transaction Solution – high level time stamps

Our offering

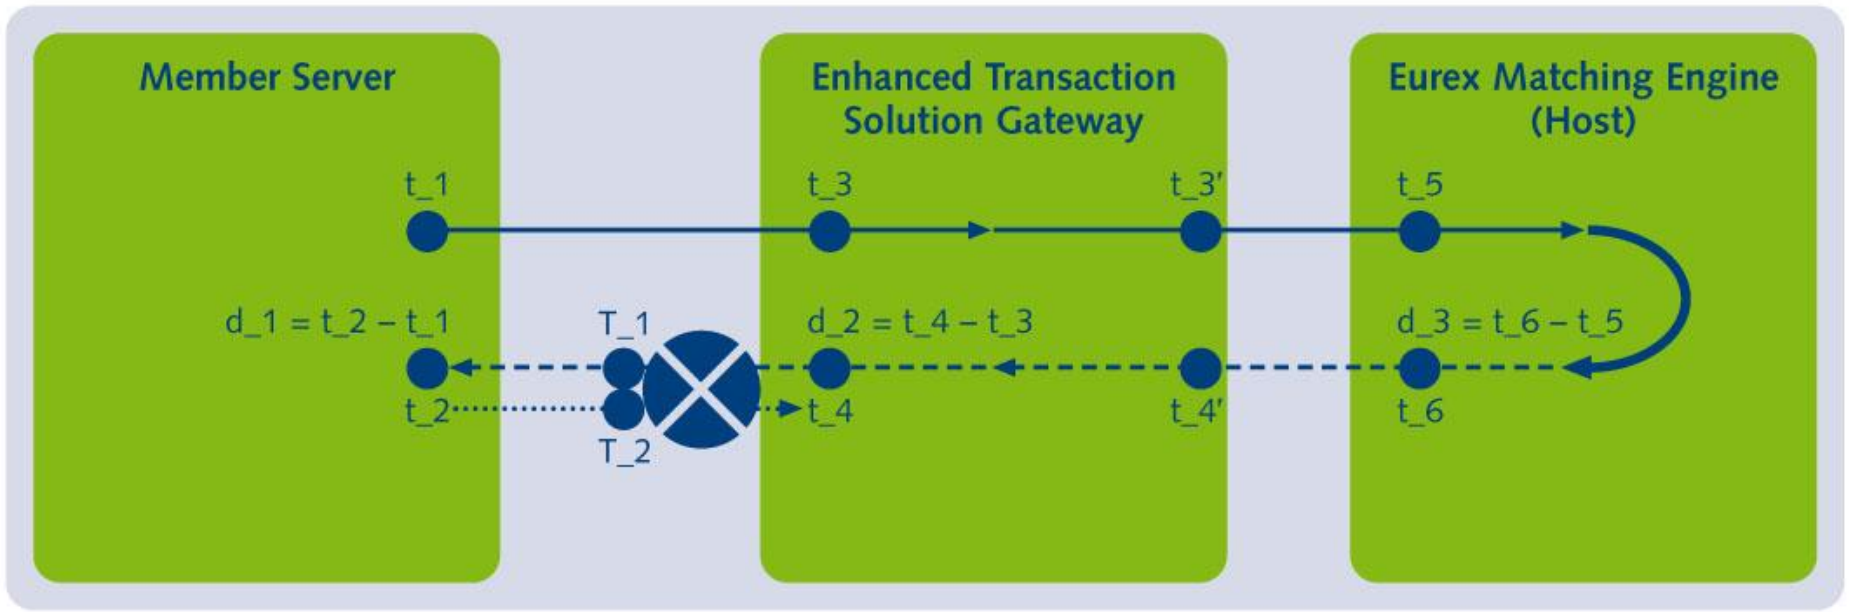

- Daily averages (per Enhanced Transaction Solution session) are available for d_2 & d_3 (Service Point at www.eurexchange.com -> technical services-> monitoring-> Eurex -> Roundtrip monitor Eurex ETS).
- Daily averages (per Enhanced Transaction Solution session) of the TCP/IP round trip times T_2-T_1 are available as well (Service Point at www.eurexchange.com -> technical services-> monitoring-> Eurex -> TCP/IP Roundtrip Statistics Eurex ETS).
- Please note that the timestamps t_3' and t_4' are not available for trading participants.

Enhanced Transaction Solution – processing times

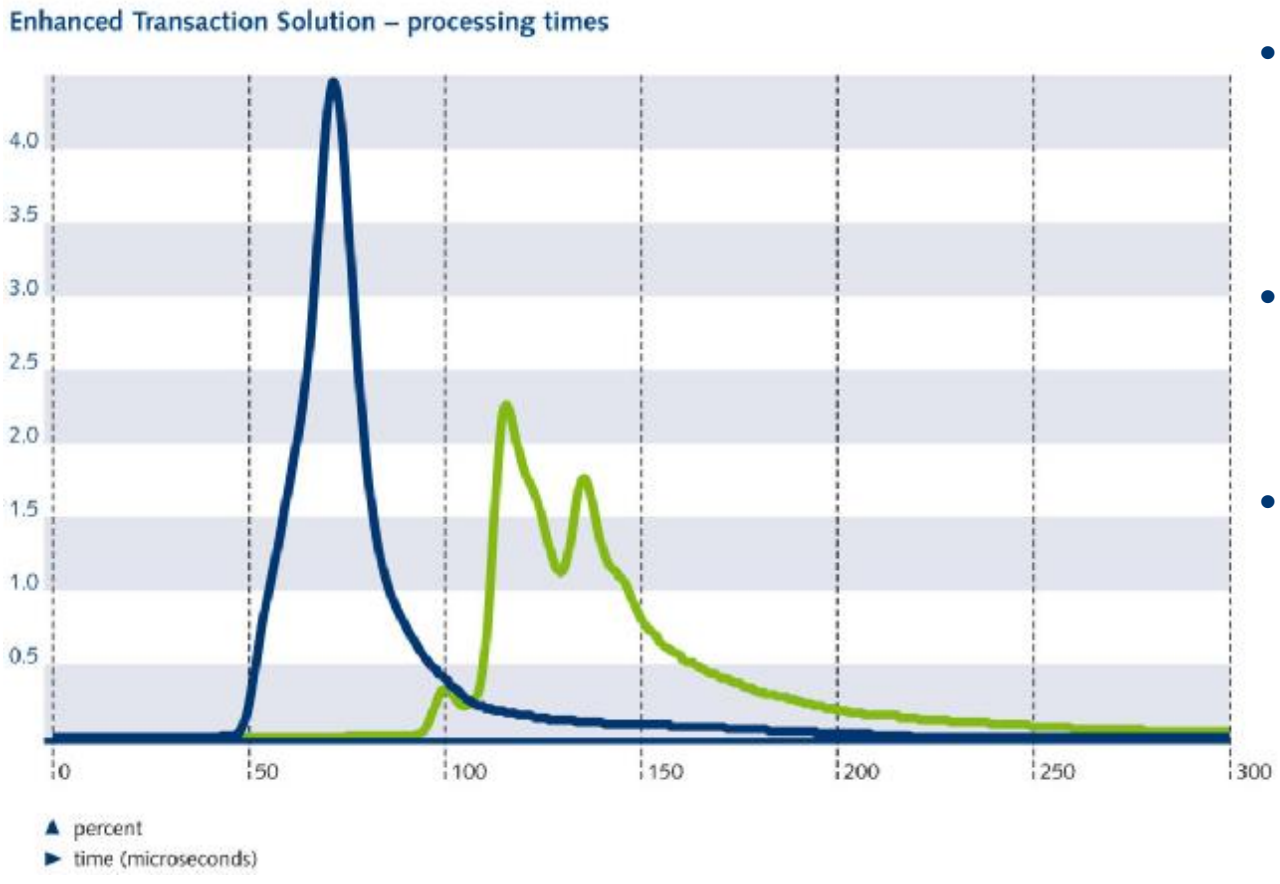

- Both graphs are histograms as of January 13, 2012 of all order/quote transactions entered via gateways located in Equinix using a bucket size of 1usec.
- Here we have plotted:
  - blue: (t_3'-t_3) + (t_4-t_4')
  - green: (t_4'-t_3') - (t_6-t_5)
- This means that the blue line is the application space processing time on the gateways whereas the green line represents the transport time between the gateways and matching engines.

- The diagram shows that the bulk of the processing time that is not spent on the matching engine is actually spent for the transport of messages from the gateways to the matching engine and back. Furthermore, one can see clearly the impact of the OptiGatewayLocID on the round trip times between the Eurex Exchange gateways in Equinix and our matching engines.

Time synchronization: Equinix gateways and matching engines

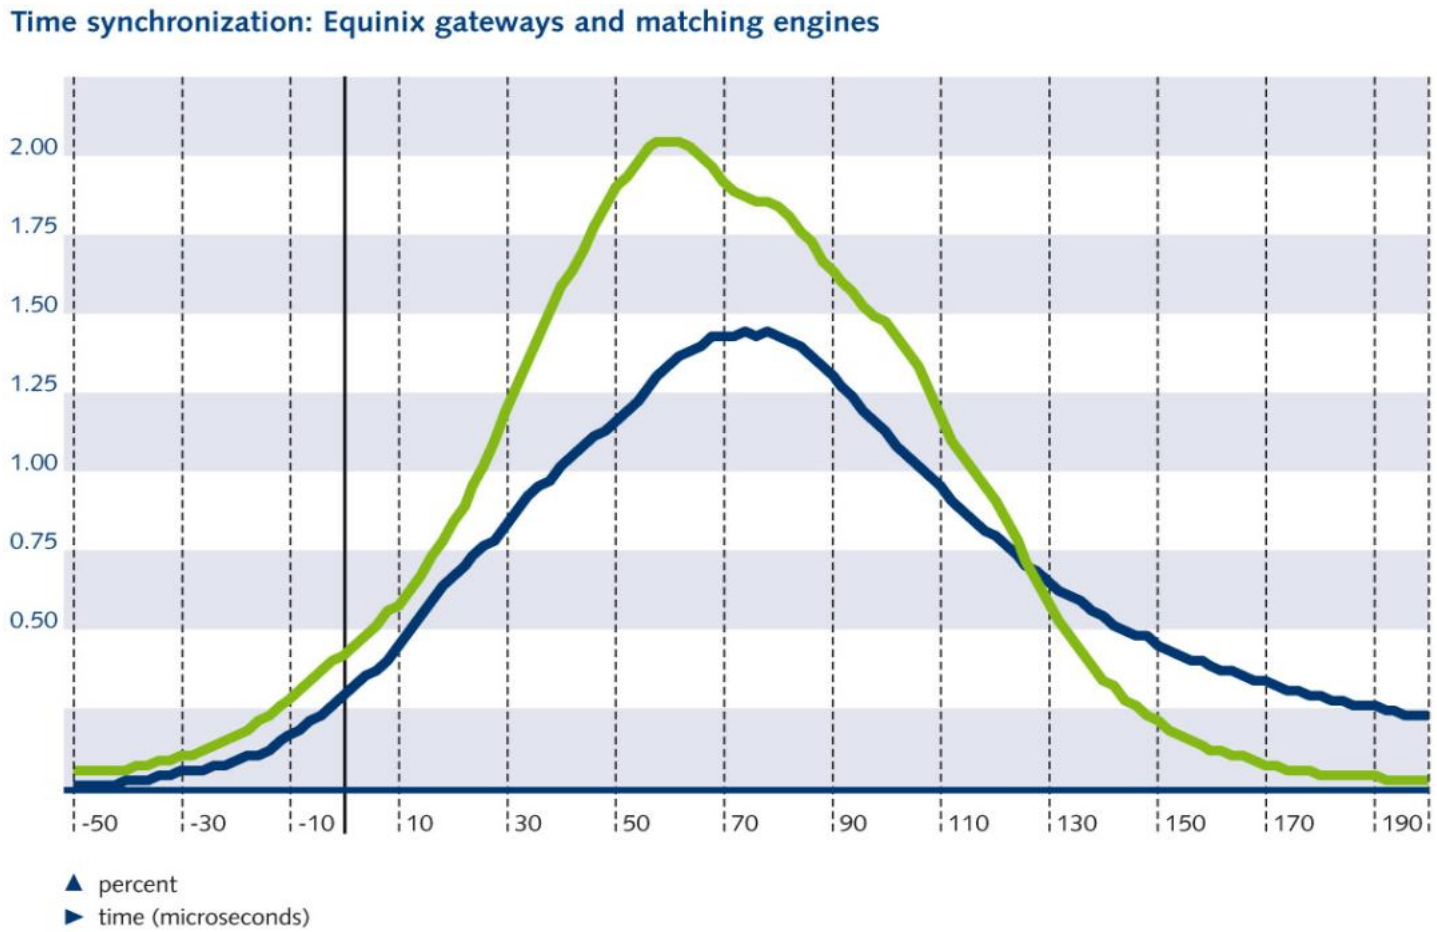

- Both graphs are histograms as of January 13, 2012 of all order/quote transactions entered via gateways located in Equinix using a bucket size of 2usec.
- We have applied a shift of the time of 480usec on the matching engine in the past.
- Here we have plotted:
  - green: (t_4'-t_6)
  - blue: (t_5-t_3')

- The diagram shows that shifting the matching engine clocks backwards by 480usec compared to the gateway clocks delivers almost no "negative" transport times. Applying this constant shift it seems that for about 99 percent of the transactions the different clocks seem to be synchronized to a level of about +/-100usecs.

Enhanced Transaction Solution – high level time stamps

Definition

- t_1,t_2: can be taken by the application when the request/response is read from/written to the socket.
- t_3,t_4: taken by the Enhanced Transaction Solution gateway when request/ response is read from/written to the socket on the participant´s side of the gateway; contained in (private) Enhanced Transaction Solution response.
- t_3',t_4': taken by the Enhanced Transaction Solution gateway when request/ response is read from/written to the socket to/from the matching engine. This data is not available to participants.
- t_5, t_6: taken by the matching engine when request/response is read from/written to the socket; contained in (private) Enhanced Transaction Solution response (note that consecutive t_5 timestamps will differ by at least 18usec).
- T_1: time when a TCP packet sent by the Eurex gateway to the participant server passes a switch just next to the Eurex gateway.
- T_2: time when the TCP acknowledge form a participant server of a TCP packet sent by the Eurex gateway to the participant server passes a switch just next to the Eurex gateway.
- A definition which Enhanced Transaction Solution/Enhanced Broadcast Solution fields correspond to the above timestamps can be found in the appendix.

Enhanced Transaction Solution – detailed futures data

Our transparency

- For the top 10 futures products daily statistics about the matching engine processing times as well as for the Enhanced Transaction Solution gateway processing times are provided via Eurex's website member section.

Sample: August 8, 2011; 08:00–22:00 (CET)

Product	Product ID	Matching Engine Round-trip times (t_6 – t_5)			Enhanced Transaction Solution Round-trip times (t_4 – t_3)		
		Average	Median	99th percent	Average	Median	99th percent
EURO STOXX 50® Index Futures	FESX	0.759	0.360	5.156	1.305	0.766	7.078
DAX® Futures	FDAX®	0.622	0.341	4.317	1.149	0.766	5.990
SMI® Futures	FSMI	0.608	0.328	4.473	1.031	0.734	5.202
STOXX Europe 50® Index Futures	FSTX	0.490	0.375	2.932	0.945	0.817	3.450
MDAX® Futures	F2MX	0.565	0.410	3.809	0.960	0.808	4.317
Euro-Buxl® Futures	FGBX	0.578	0.410	3.577	1.027	0.861	4.166
Euro-Bund Futures	FGBL	0.987	0.476	6.657	1.606	0.936	8.729
Euro-Bobl Futures	FGBM	0.776	0.426	5.682	1.383	0.908	7.658
Euro-Schatz Futures	FGBS	0.436	0.276	2.008	0.988	0.775	3.642
Long-Term Euro-BTP Futures	FBTP	0.460	0.405	1.453	0.906	0.852	1.990

- Please note that the statistical data set is the set of all non-error, non-persistent order and quote related transactions for the respective Eurex Exchange product sent via a Enhanced Transaction Solution session on gateways in the Equinix data center.

Agenda

- Achievements
- Topology
- Inside the matching engine
- Infrastructure
- Outlook
- Appendix

Overview

Inside the matching engine – overview

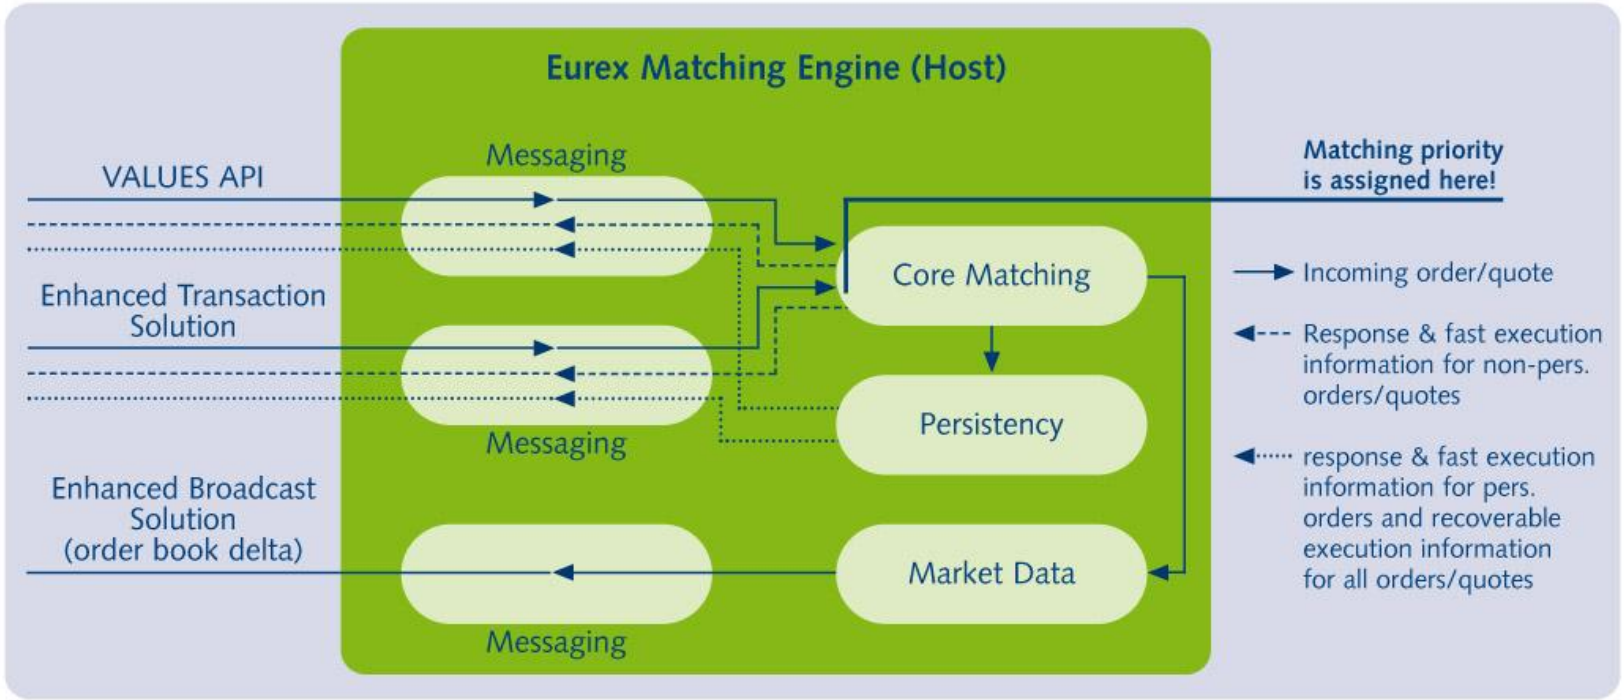

- Orders/quotes entered via VALUES API go through a MISS and a so-called communication server (not streamlined for low latency).
- Orders/quotes entered via the Enhanced Transaction Solution go through a gateway (streamlined for low latency).
- In order to get the best possible latency you have to use proximity access to the Eurex gateways as the (Eurex internal) network to connect to these gateways has been streamlined compared to network connectivity from elsewhere.

Details (I)

- Orders/quotes entered for a specific product are sent by the gateway/communication server to the respective Eurex matching engine.
- Here they are forwarded by two messaging components (for orders/quotes entered via VALUES API and via Enhanced Transaction Solution) to the core matching component.
- The matching priority is assigned when the orders/quotes are handed over to the core matching component.
- Statistically most “delays” in the Eurex system happen here during high load situations.

Details (II)

- The core matching component works as follows:
  - when an order/quote arrives, it is functionally processed (e.g. put in the book).
  - resulting messages for non-persistent orders/quotes are sent out in the following order:
    - direct response to the order/quote entered
    - fast execution information for booked orders /quotes and fast execution information for the aggressing order/quote
  - handover of data to the market data component which sends out an order book delta message as well as trade messages in case of a match.
  - resulting messages for persistent orders and executions are sent to the persistency layer.
  - after persistence, the resulting messages for persistent orders and executions are sent out. These are: the response & fast execution information for persistent orders as well as the recoverable execution information for all orders.

Details (III)

- In case that during this phase several new orders/quotes transactions arrive at the core matching component:
  - The core matching component will batch these orders/quotes and treat them in one atomic unit (direct responses and fast execution information will be sent out one-by-one for non-persistent order/quotes as soon as the core matching component has worked out the details).
  - At the end of the atomic unit all trade prices and the final order book will be handed over to the market data component and published.
  - Data will be handed over to the persistency layer and the resulting messages for persistent orders and executions are sent out. These are: the response & fast execution information for persistent orders as well as the recoverable execution information for all orders.

Time stamps (I)

Inside the matching engine – time stamps

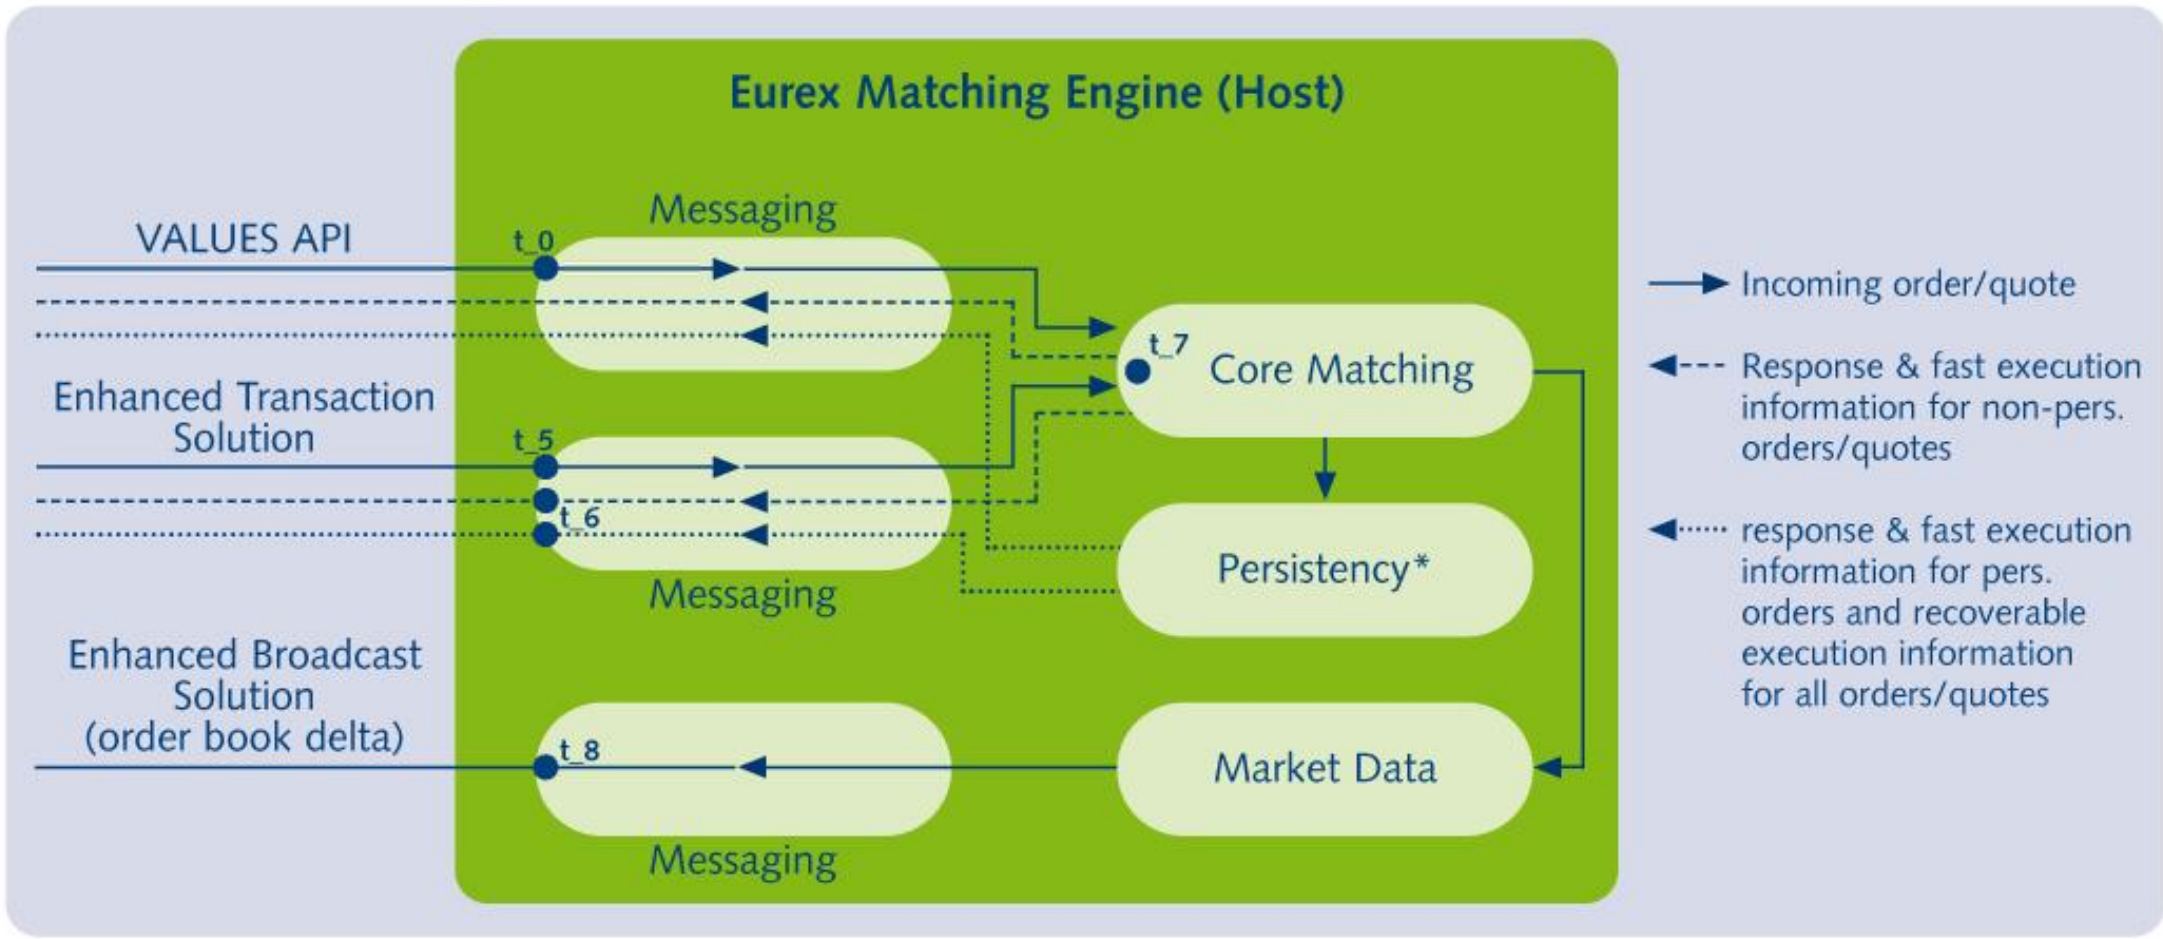

* Additional time stamps are available in messages from the persistency layer.

Time stamps (II)

- t_0, t_5: taken when an order/quote is read from the socket; contained in the Enhanced Broadcast Solution trade message in case of a match (note that consecutive t_0 or t_5 timestamps will differ by 18-19 usec at least).
- t_6: taken when the response of an order/quotes transaction or execution information is written to the socket; contained in the (private) Enhanced Transaction Solution response.
- t_7: taken when the order is functionally processed (e.g. put in book or matched); available in the (private) Enhanced Transaction Solution response and Enhanced Broadcast Solution order book delta and Enhanced Broadcast Solution trade message in case the order matches.
- t_8: taken just before a market data UDP datagram is written to the socket.
- For every Enhanced Broadcast Solution datagram that contains an order book incremental message, Eurex publishes a payload indicator in the version information message. The value of the indicator is the time difference t_8-t_5 (or t_8-t_0). Here t_5 (or t_0) is the time when the "oldest" order which is contained in this market data update reached the Eurex matching engine.
- A definition which Enhanced Transaction Solution/Enhanced Broadcast Solution fields correspond to the above timestamps can be found in the appendix.

Time stamps (III)

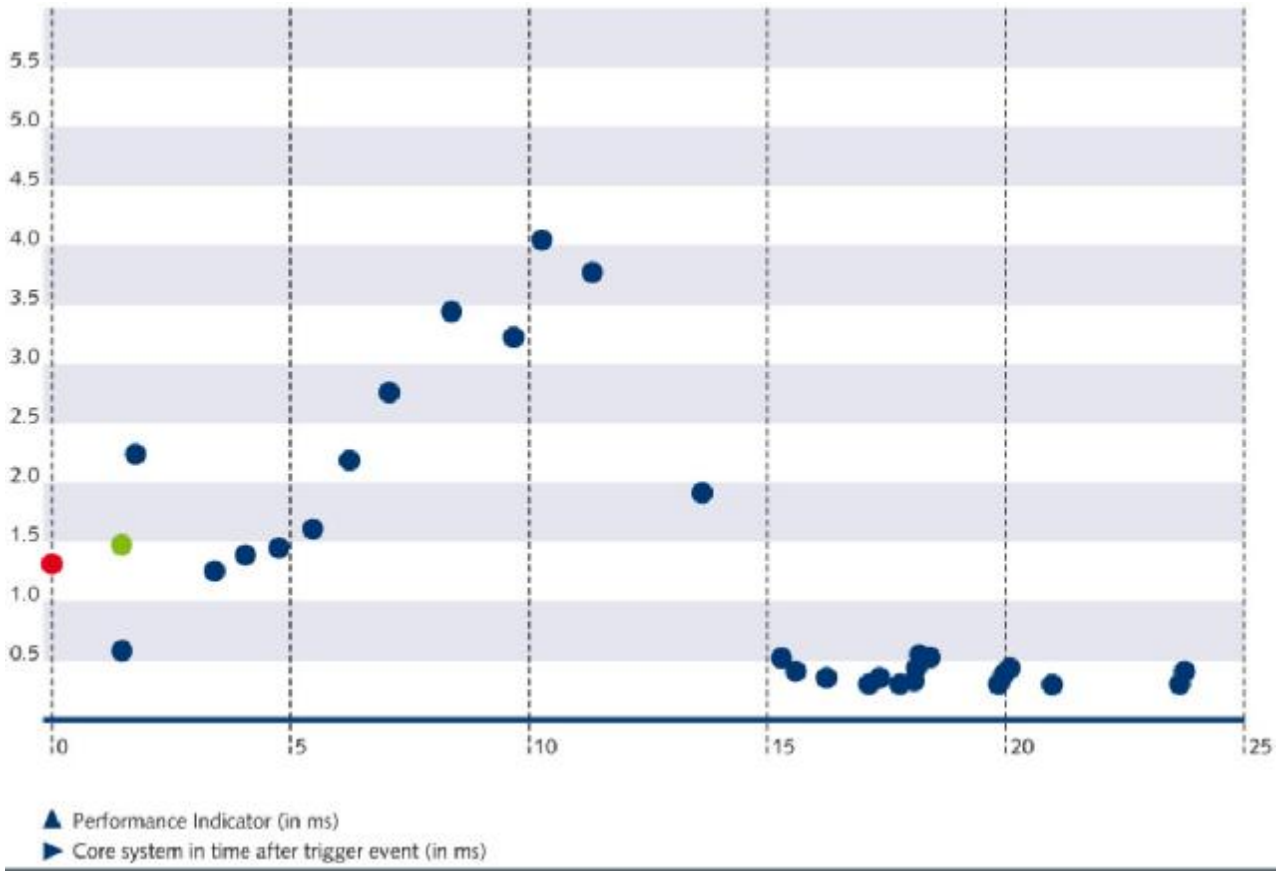

- This is an explicit example of the usefulness of time stamps, specifically the Performance Indicator (PI).
- The PI is defined by the time it took the host taking a order/quote from the wire (core system in-time), processing it, and putting market data back on the wire (Enhanced Broadcast Solution send time*).
- The graph depicts the performance development following a large trade in EURO STOXX 50® Index Futures (FESX) on March 28, 2012.

- We plotted the PI (vertical) against the host in-time of the oldest order/quote in the respective Enhanced Broadcast Solution order book delta (host in-time equals Enhanced Broadcast Solution send time minus the PI).
- The red dot represents the Enhanced Broadcast Solution packet that contains the large trade, which triggered a burst of incoming orders.
- The green dot represents the Enhanced Broadcast Solution order book delta that is caused by the "winning" order, which resulted in a trade. The subsequent reduced performance lasted about 15 ms.

** Whenever an Enhanced Broadcast Solution order book update covers multiple orders/quotes, the PI relates to the oldest order/quote.*

Agenda

- Achievements
- Topology
- Inside the matching engine
- Infrastructure
- Outlook
- Appendix

What you need to be fast…

A few recommendations to achieve the lowest possible latency:

- Use the Equinix co-location facility to be close to the Eurex Exchange system.
- Try the two different Enhanced Transaction Solution gateways assigned to each session to see which delivers the better performance for your strategy (try it out and compare your time stamps as well as p&l for different days).
- Use the Enhanced Transaction Solution fast info notifications for trading decisions although executions reported by such messages are only indicative (and not legally binding).
- Recoverable messages need to be processed to create safety. Therefore, we recommend to use either a FIX trade capture drop copy or the Enhanced Confirmation Solution (using either trade confirmations or matching events) to confirm the fast info notifications*.
- Use 10 GE (cross-) connections for the Enhanced Broadcast Solution and Enhanced Transaction Solution.
- Use a state-of-the-art switches and only have at most one hop between the exchange network and your server.
- Use good NIC cards or TCP/IP acceleration, e.g. a Linux kernel-by-pass library.

** Please see Eurex circular 268/2011 for more details and best practices in handling of the "optimistic" Enhanced Transaction Solution information notifications.*

Agenda

- Achievements
- Topology
- Inside the matching engine
- Infrastructure
- Outlook
- Appendix

Outlook

Eurex Exchange has announced the launch of its new trading architecture for December 2012.

The new system will deliver:
- state-of-the-art technology which is built on the pillars performance, efficiency, capacity, and reliability,
- minimize latency, maximizing throughput and allowing for greater flexibility while
- maintaining high standards of reliability.

With the new trading architecture available, Eurex Exchange will cease to use the currently used MISS infrastructure and VALUES API interface for trading.

Eurex Exchange offers co-location clients with at least two 10 GE cross connects one additional 1 GE fiber cross connect delivering access to a high quality, GPS synchronized time service in form of NTP and PTP signals for free*.

An update of this presentation for the new trading architecture is provided prior to the launch of the new system on December 3, 2012. For more information about the new trading architecture please visit www.eurexchange.com/nta

Finally, Eurex Exchange will continue to investigate possibilities of extending the transparency with respect to latency figures - in case you have any suggestion – please get in touch with us!

**Please see the following link for more details:*
http://deutsche-boerse.com/it/dispatch/en/kir/gdb_navigation/technology/30_Access_Products/20_ExServes/80_ExServes_Time_Service

Thank you for your attention!

Agenda

- Achievements
- Eurex Topology
- Inside the matching engine
- Infrastructure
- Outlook
- Appendix

Timestamp Reference

The timestamps t_3,…,t_8 are available via the following fields:

- t_3: gatewayReqInTimeStamp in the Enhanced Transaction Solution TM Response Notification
- t_4: gatewayRespOutTimeStamp in the Enhanced Transaction Solution TM Response Notification
- t_5: coreSystemInTimeStamp in the Enhanced Transaction Solution TM Response Notification
- t_6: coreSystemOutTimeStamp in the Enhanced Transaction Solution TM Response Notification
- t_7: lastEventTrnId in the Enhanced Transaction Solution Enter/Modify/Delete Order/Quote/Mass Quote Response and
  hiResTimestamp in the Enhanced Broadcast Solution Order Book Incremental or
  entryTim in the Enhanced Broadcast Solution Trade Information Message
- t_0, t_5: aggressorEntryTim in the Enhanced Broadcast Solution Trade Information Message
- t_8: timeStamp in the Enhanced Broadcast Solution UDP packet header
- (t_8-t_5) , (t_8-t_0):
  performanceIndicator in the Enhanced Broadcast Solution UDP packet header