

Open Day – Deutsche Börse IT conference 2011

Latency in the Deutsche Börse network

Michael Neuerburg / Stephan Hoppe

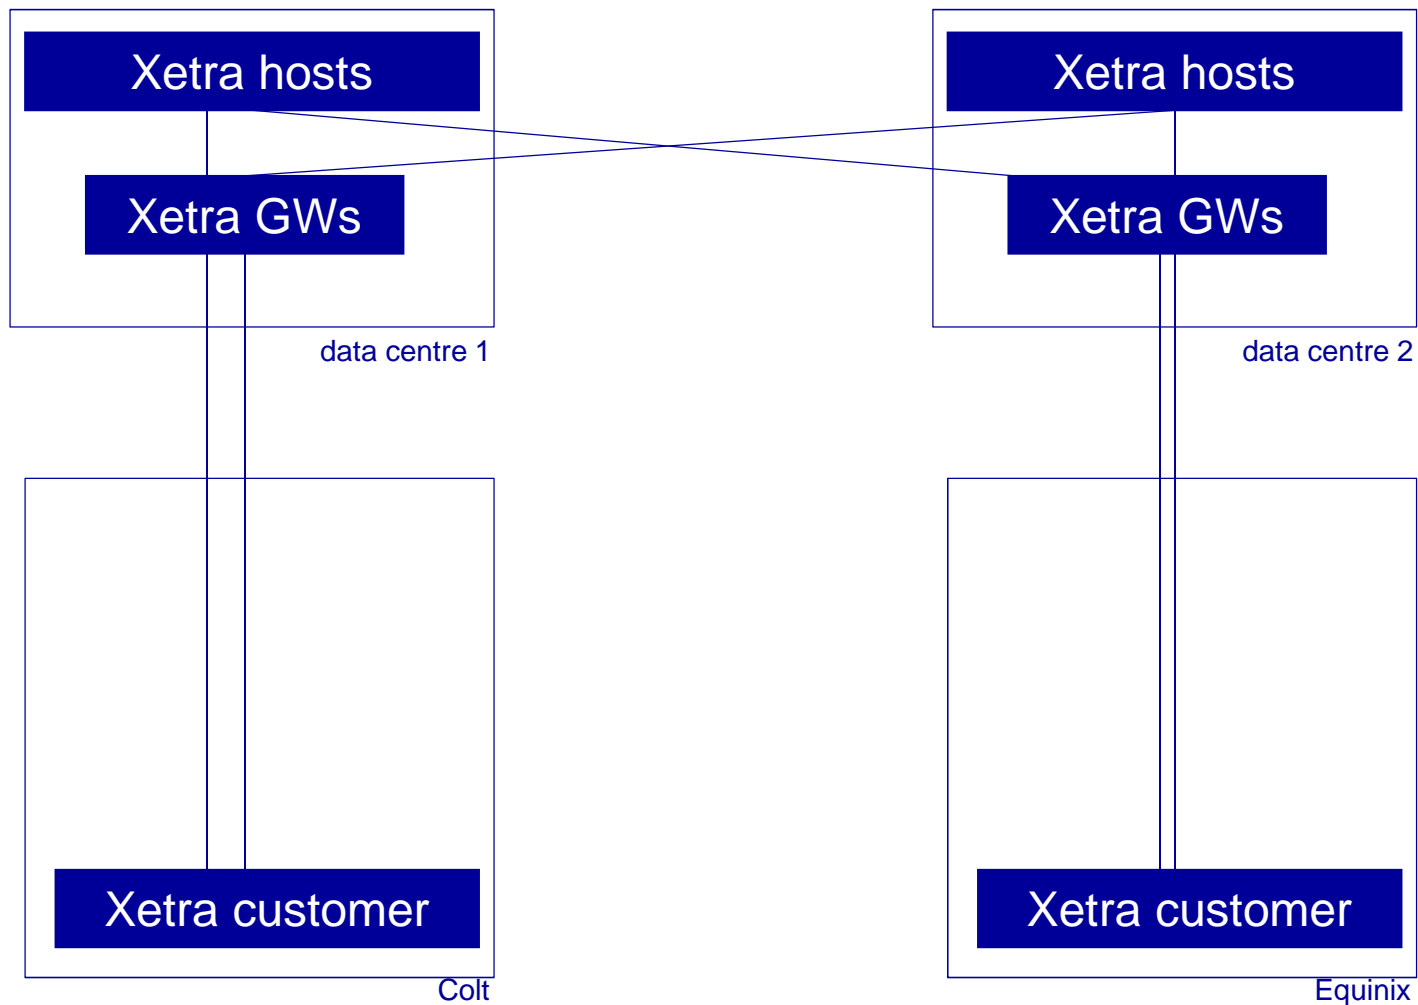
Eschborn, 5 October 2011



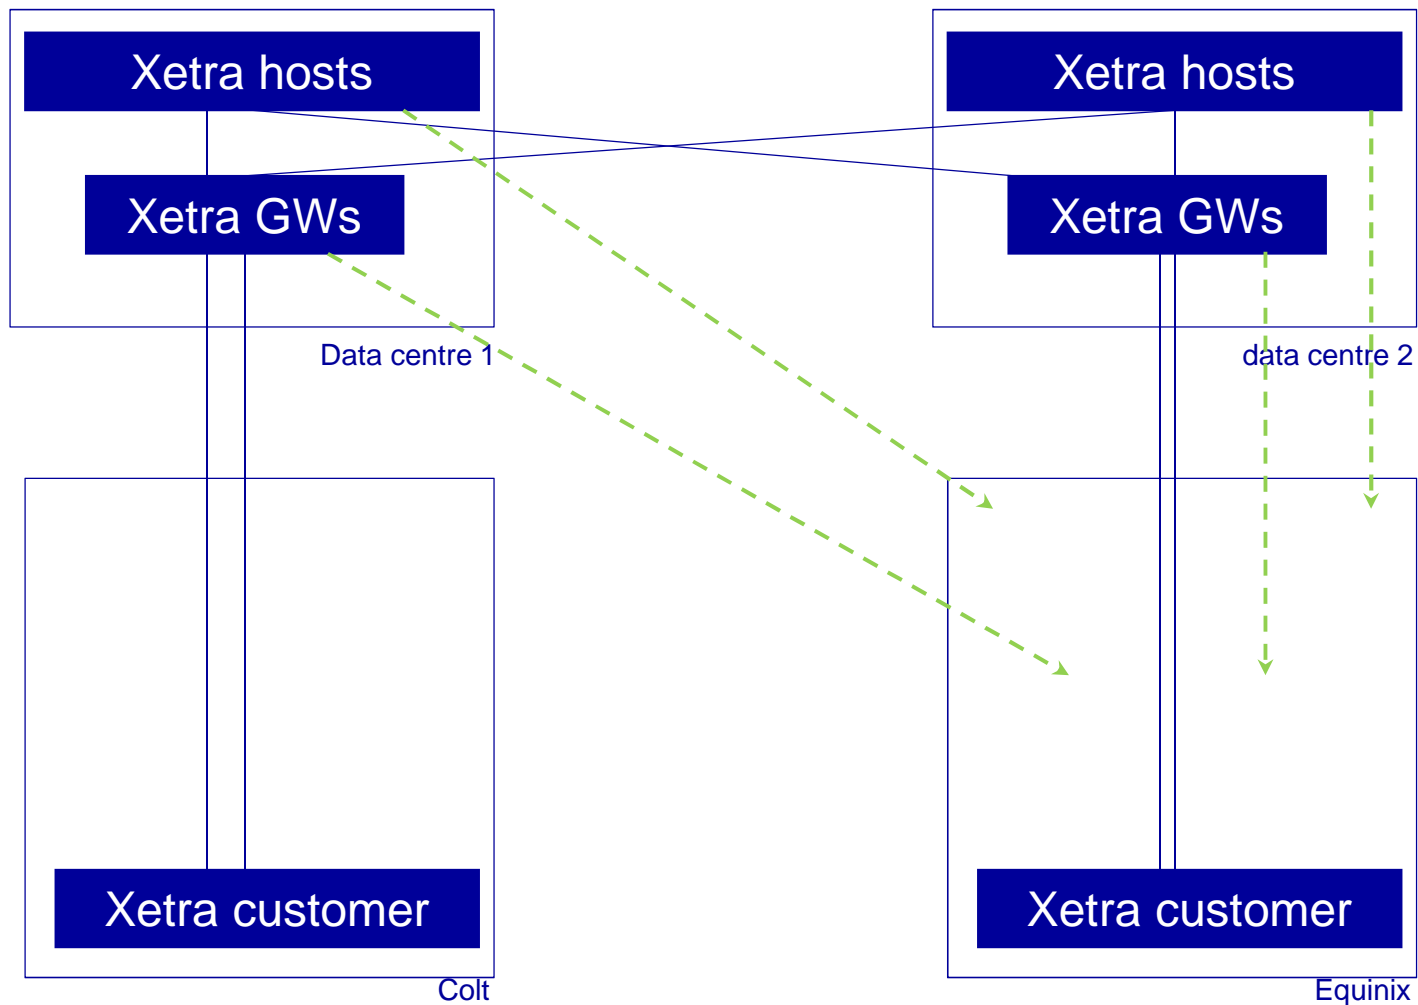
Agenda

- n Co-location for Xetra[®]
 - n Towards Co-location for Eurex[®]
 - n Redesign of the backbone network
 - n Disaster recovery in the new constellation
 - n Outlook
 - n (Appendix: Timestamps in Eurex Enhanced Transaction Solution / Enhanced Broadcast Solution)
- } results

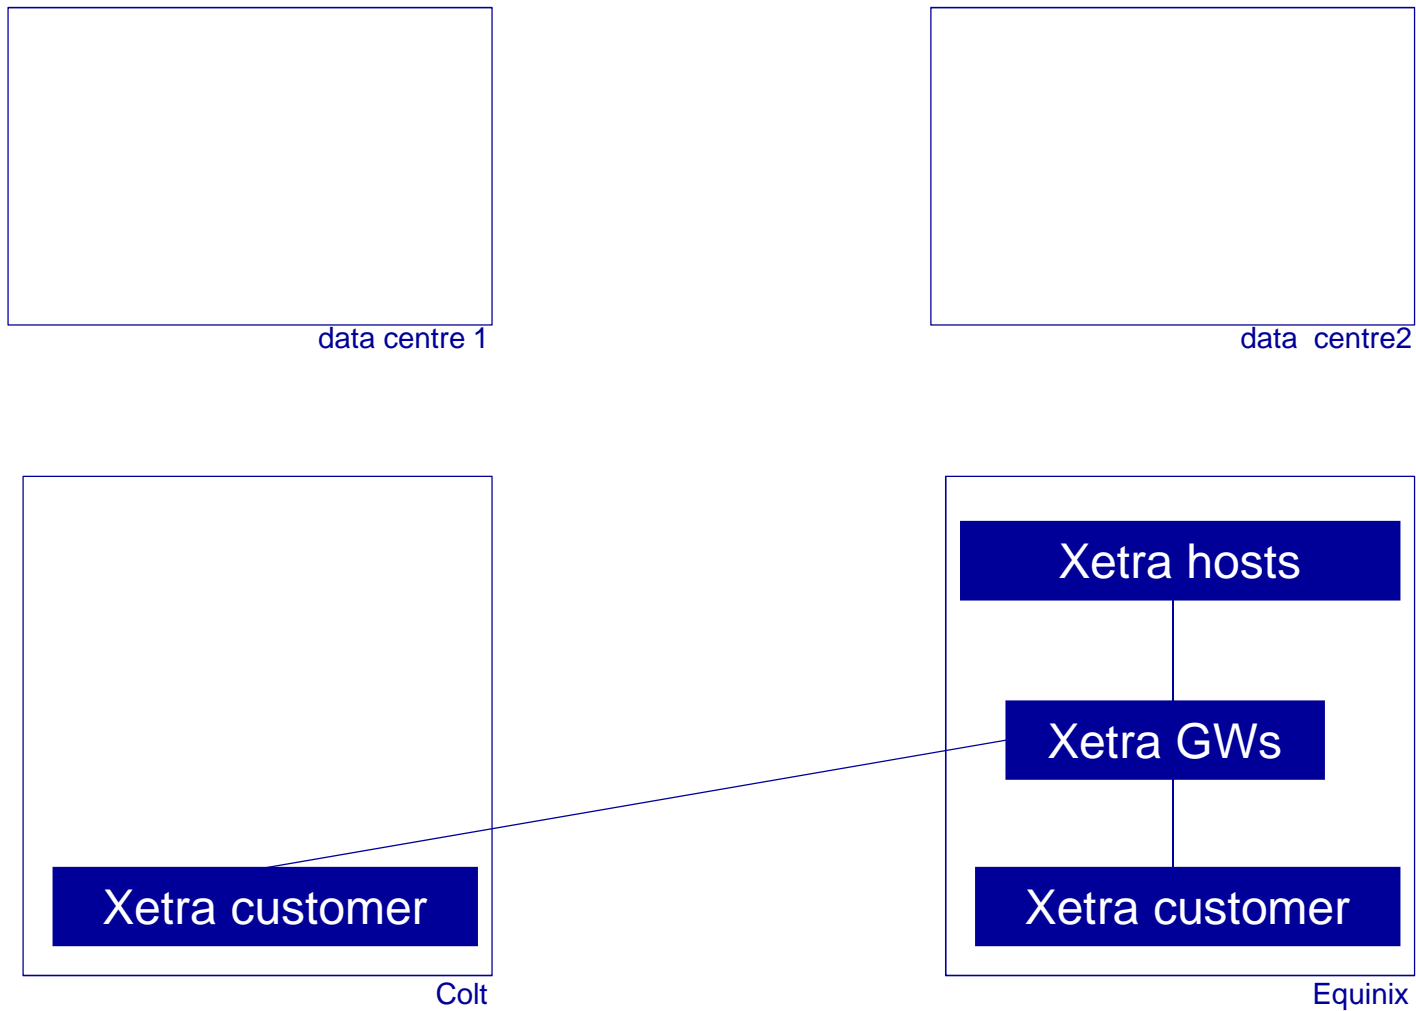
Xetra[®]: The move to co-location



Xetra[®]: The move to co-location



Xetra[®]: The move to co-location



Xetra[®]: The move to co-location

Expectations and results for request/response traffic

The physical distance between co-location customer and gateway has been drastically reduced

- Therefore the significance of the published Enhanced Transaction Solution latency numbers is now less important than it had been. However, it still is a means to detect whether there are serious problems.

31 May

Location	TCP/IP Packets	Avg. Latency [ms]	Median Latency [ms]
PRO	115786	0.179	0.176
PRO	115647	0.18	0.176
PRO	110821	0.177	0.176
PRO	111834	0.221	0.219
PRO	109973	0.177	0.176
PRO	117259	0.22	0.219
PRO	116119	0.177	0.176

1 Aug

Location	TCP/IP Packets	Avg. Latency [ms]	Median Latency [ms]
PRO	40982	0.051	0.049
PRO	41113	0.052	0.049
PRO	35187	0.05	0.049
PRO	35196	0.05	0.049
PRO	23823	0.051	0.048
PRO	40680	0.053	0.049
PRO	38821	0.046	0.044

(no problems here!)

Xetra[®]: The move to co-location

Expectations and results for request/response traffic

Due to replacement of the hardware for the gateway with a current state-of-the-art model

- n we can observe an improvement in the gateway internal processing time of about 60 μ s

As gateway and host are now in the same building, the physical distance has been reduced significantly

- n This leads to an improvement in the network latency between host and gateway (depending on former product location) of about 75 μ s and 140 μ s respectively

As the hardware for the hosts was replaced with a current state-of-the-art model

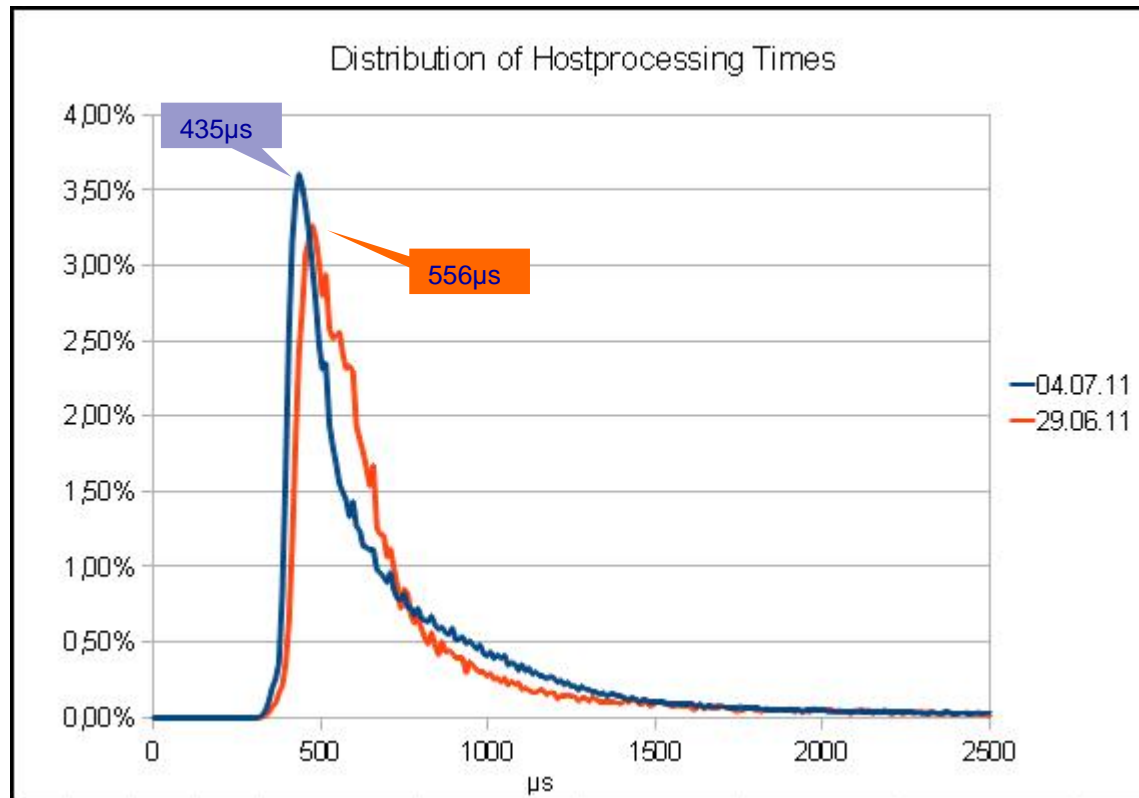
- n processing time on the backend was improved

As the whole backend cluster is now in one location

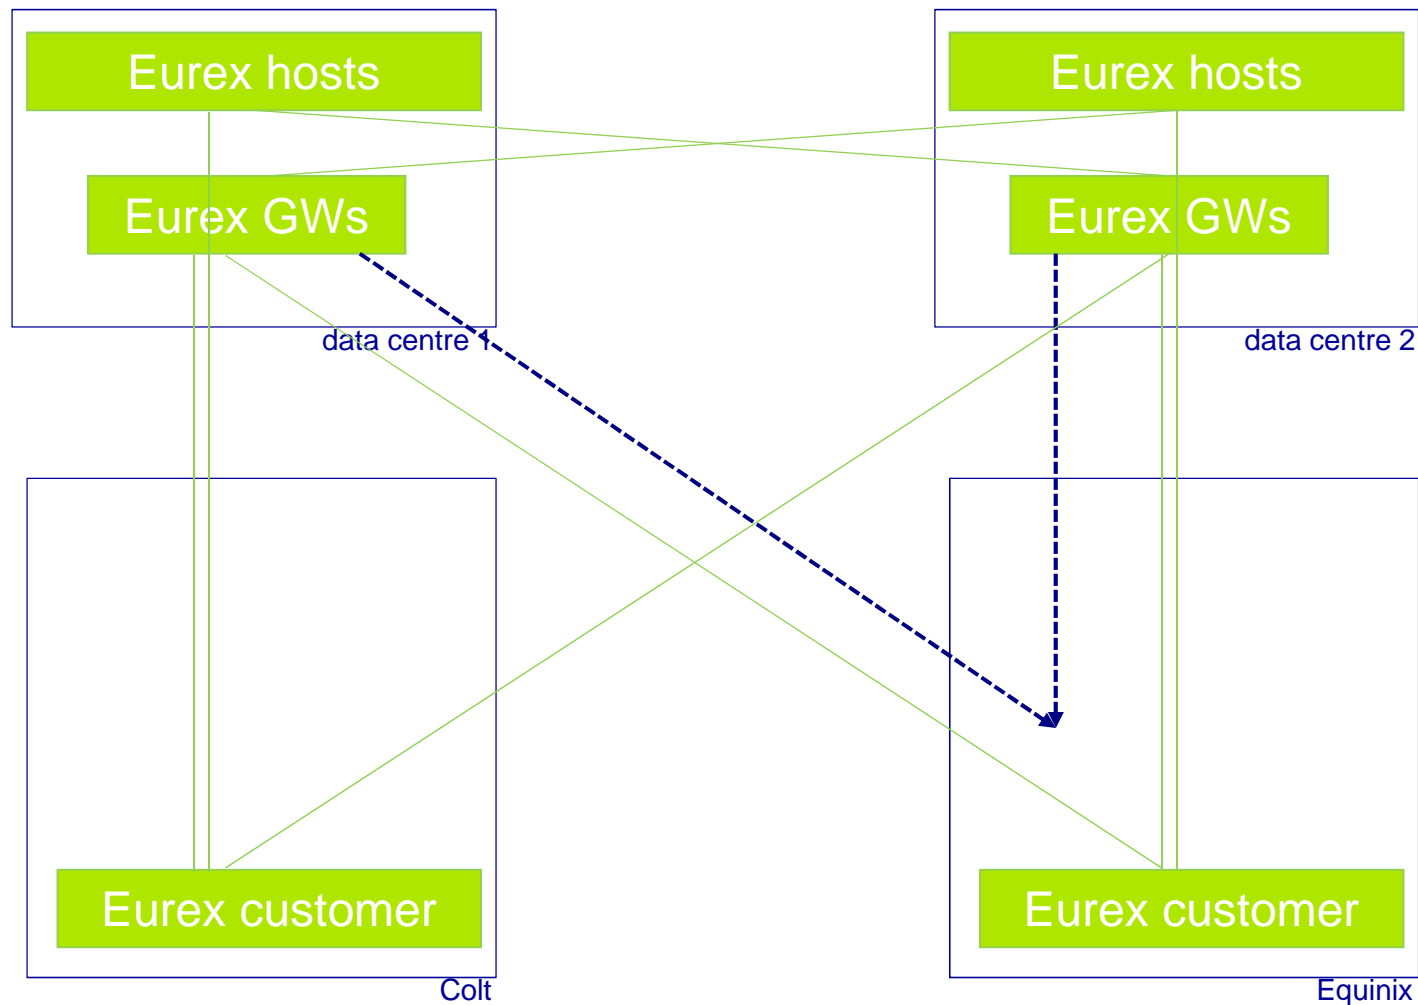
- n intra-cluster traffic is faster, leading to further improvements

Xetra[®]: The move to co-location

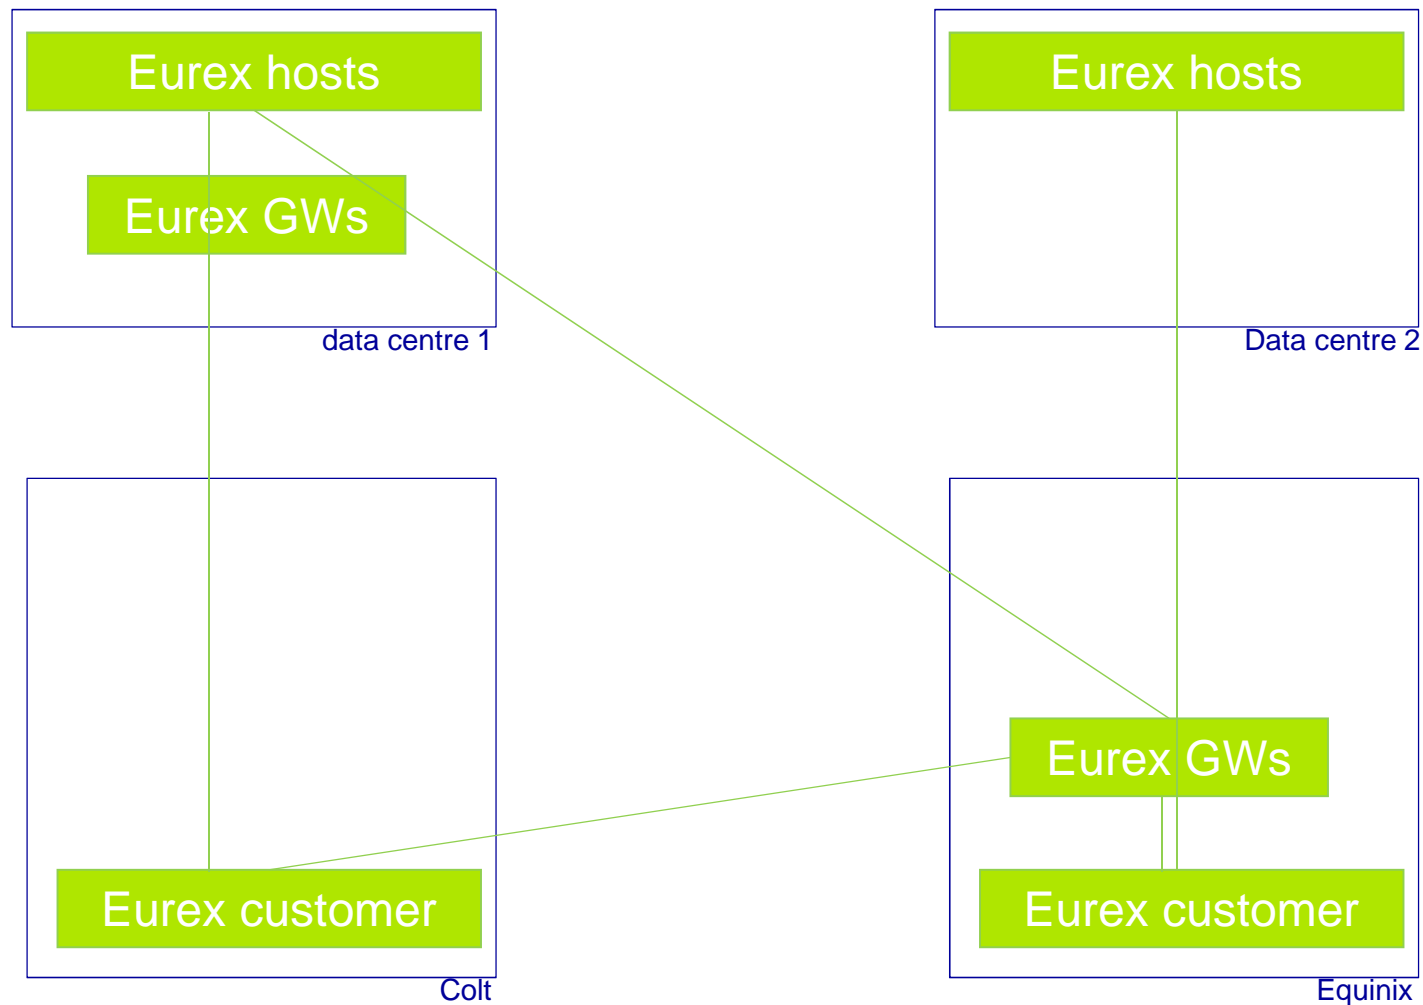
Expectations and results for request/response traffic



Eurex[®]: A first step towards co-location



Eurex[®]: A first step towards co-location



Eurex[®]: A first step towards co-location

Expectations and results: request/response traffic

Xetra[®] and Eurex are operated by the same gateway hardware. Furthermore, Equinix customers are now next to the Eurex gateways. Thus, the improvements for the gateway internal processing time, which have been observed for Xetra, apply for Eurex as well. Latency between customer and gateway is again in the 60µs range.

Even though the distance between the gateways and the Eurex hosts is now larger, the latency between customer and host should be roughly the same as before - except for the better gateway processing times.

This can be observed by evaluating the core and gateway timestamps in the Enhanced Transaction Solution. We did this for a particular customer: the results are displayed on the next slide.

Eurex[®]: A first step towards co-location

Results in Equinix for Enhanced Transaction Solution after replacement of opti=0 GW

Customer timestamps – Core timestamps (all in μ s)

Products in			31 May	16 June	Difference
Data centre 1 (opti = 0)	FDAX	Avg	806	645	161
		Median	636	510	126
	FGBL	Avg	807	641	166
		Median	572	443	129
Data centre 2 (opti = 1)	FESX	Avg	867	661	206
		Median	585	461	124
	FGBX	Avg	627	475	152
		Median	562	437	125

Xetra[®]/Eurex[®]: The move to co-location Expectations and results for broadcast traffic

Xetra:

As the paths in the broadcasting network between host and gateway are now truly symmetric...

- n the difference in arrival times for A and B stream has been reduced and reflects the difference in time between sending the first and the second packet only

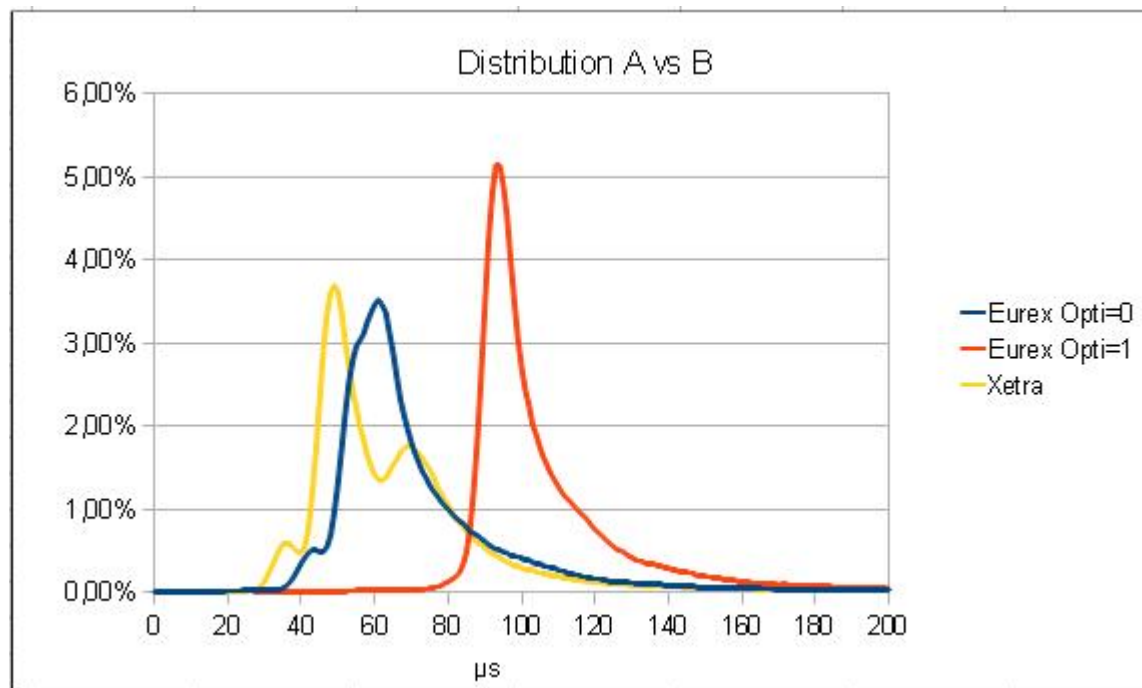
Eurex:

The products are still hosted in two datacenters.

- n Therefore there is still a notable difference between the two product classes

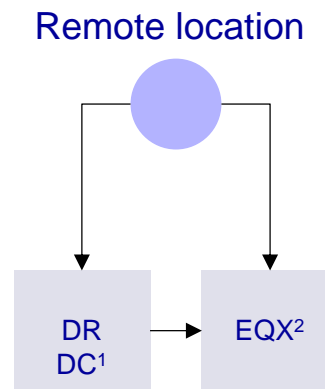
Xetra[®]/Eurex[®]: The move to co-location

Expectations and results for broadcast traffic



Redesign of backbone network (I)

- n With the host moving to the new Deutsche Börse data centre (Equinix) the low latency backbone network has been redesigned.



- n All latency optimised connections are connected to the Equinix data centre, while the second, redundant backbone is routed via the disaster recovery site.

¹ disaster recovery data centre

² Equinix

Redesign of backbone network (II); London

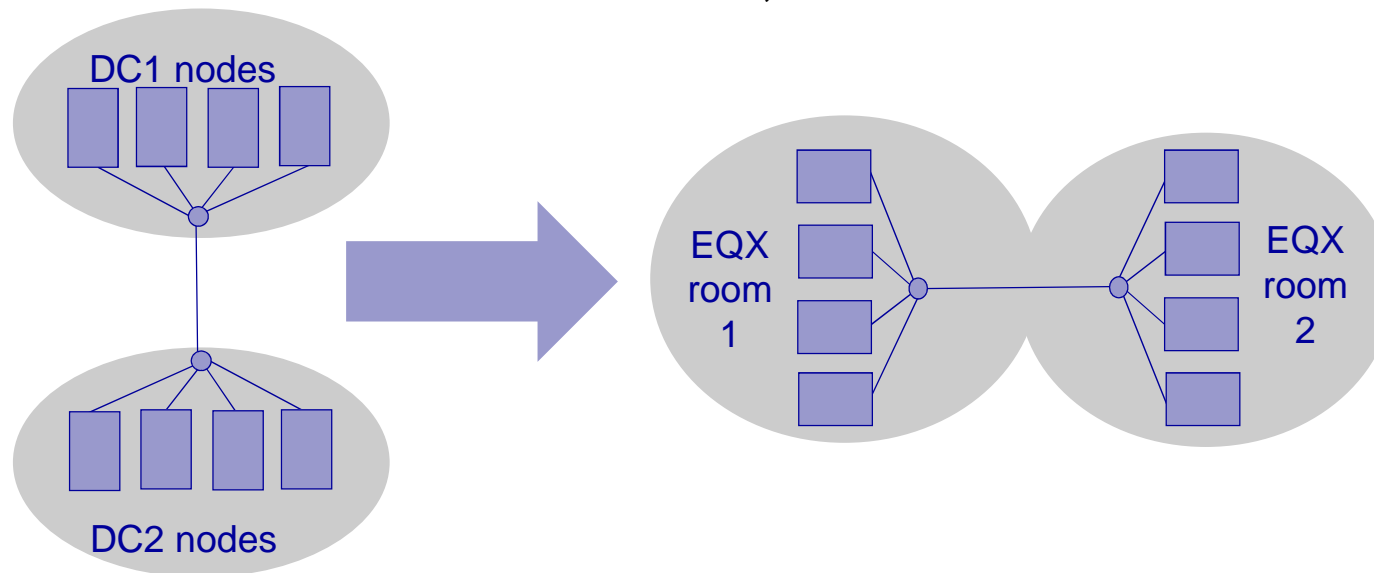
- n Using two different backbone circuits with different routing paths leads to a latency difference between both circuits.
- n Following the discussion of the last Open Day, backbones to London have been harmonised as far as possible:

Backbone London	one-way latency DR site	One-way latency production EQX
Before	4,8 ms	7,9 ms
Today	4,8 ms	4,3 ms

- n Fast backbone connects Equinix with Proximity (Telehouse North), London.

Failover and disaster recovery setup

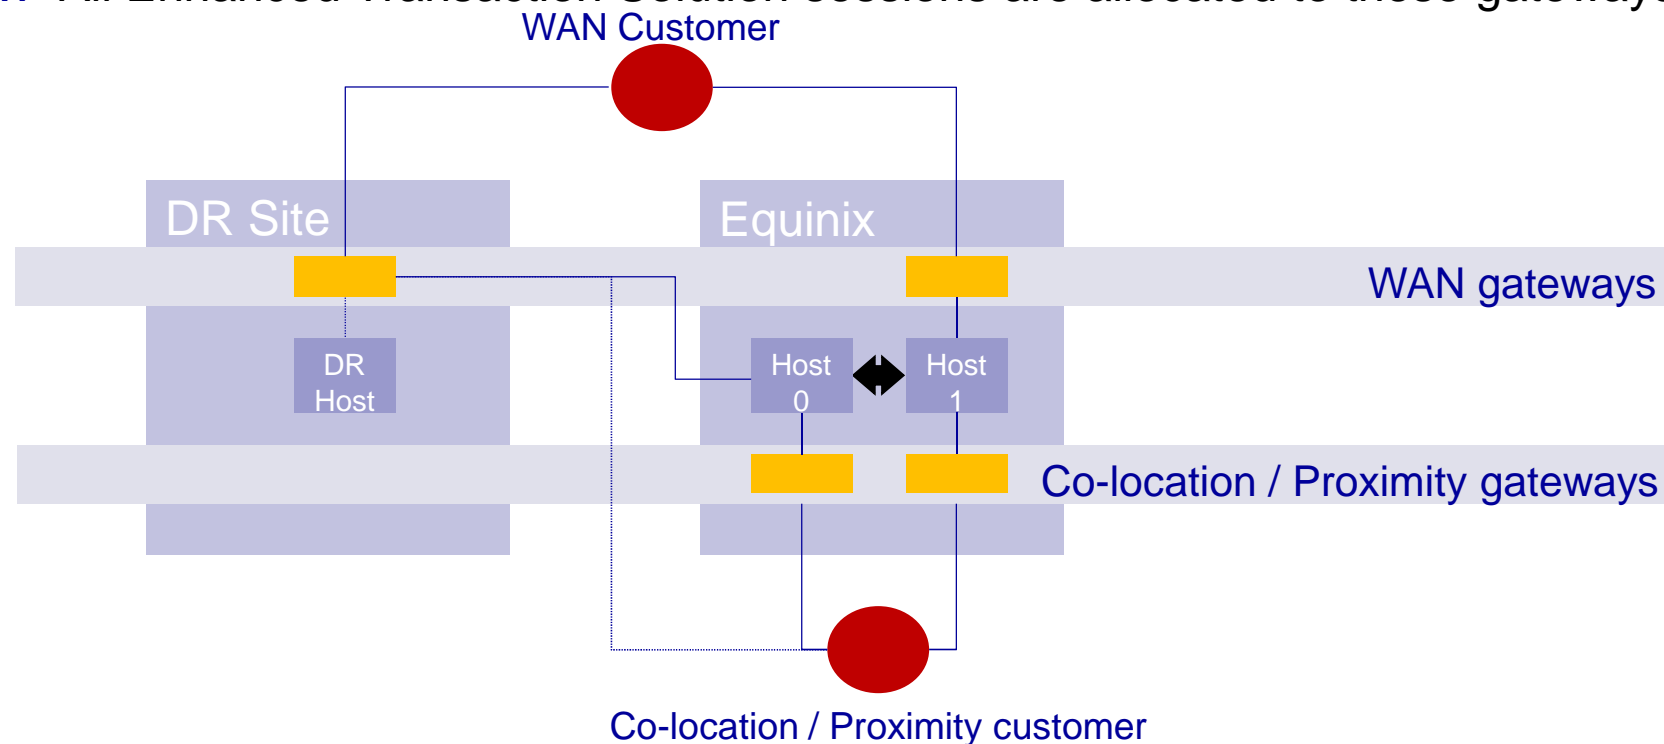
- n The Cluster of Xetra[®] and Eurex[®] is spread between two different rooms within the Equinix data centre.
- n Both rooms have been designed to be as redundant as possible to each other (with regards to power and air-conditioning).
- n In case of failure of one of these rooms, the other room will seamlessly take over.



- n In addition, a second data centre for disaster recovery is available

Example disaster recovery setup: Enhanced Transaction Solution

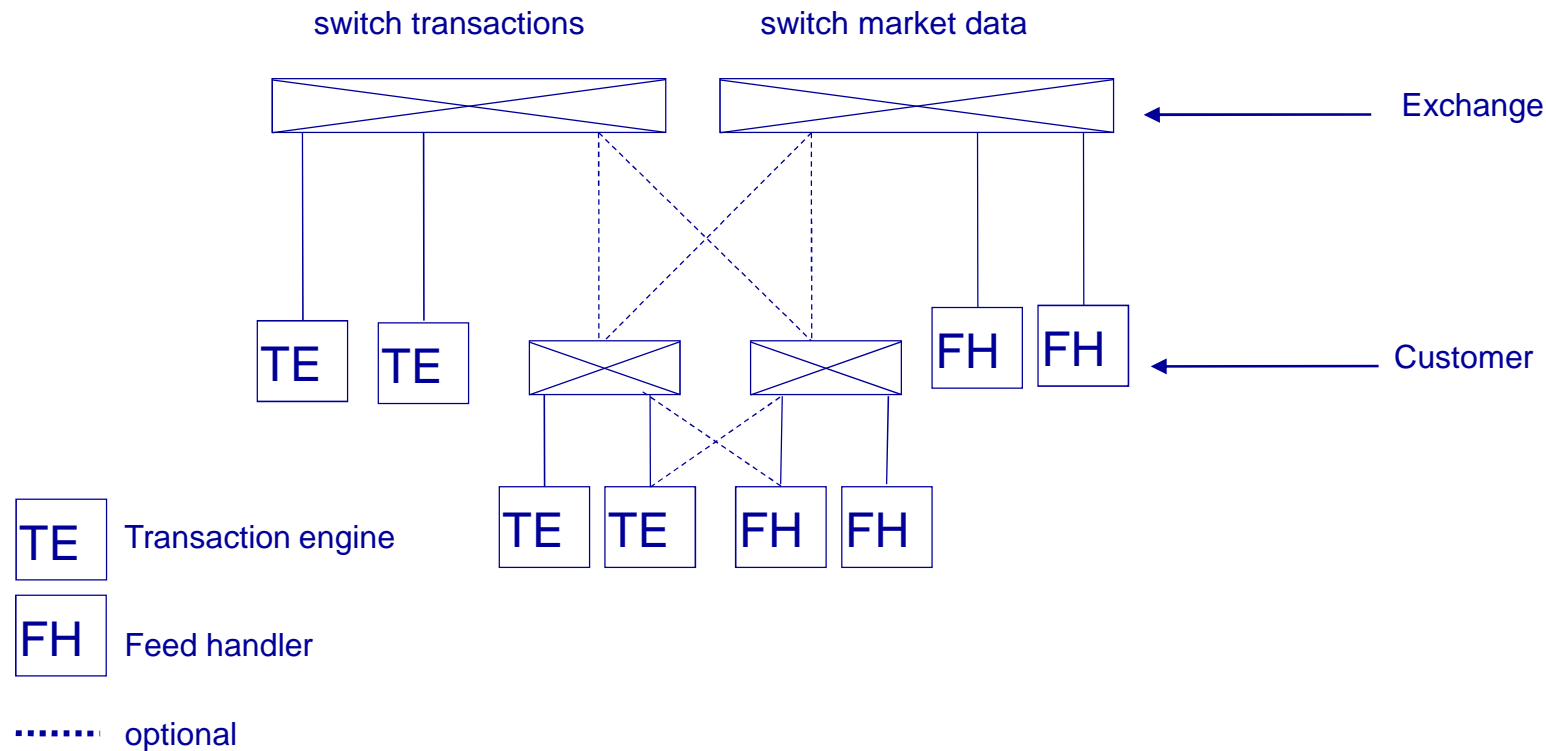
- n For Eurex[®] and Xetra[®] both host sides are located at Equinix with a disaster recovery host located in a second data centre.
- n In case of disaster recovery, production will be possible via gateways located in disaster recovery site.
- n All Enhanced Transaction Solution sessions are allocated to these gateways.



Outlook

New connection alternative for Enhanced Broadcast Solution and Enhanced Transaction Solution

- n 10 GB fibre connectivity within the Equinix data centre
- n Dedicated for Enhanced Transaction Solution and Enhanced Broadcast Solution
- n Standardised cable length from the customer rack to the host



Appendix: Eurex[®] timestamps in Enhanced Transaction and Enhanced Broadcast Solution

There is a large number of timestamp fields in our protocols.

Often the documentation is not really helpful to guess their meaning.

Furthermore, these timestamps come in various degrees of accuracy and use different baselines.

e.g. ETS.ordrCreatTim is in 100ms since midnight

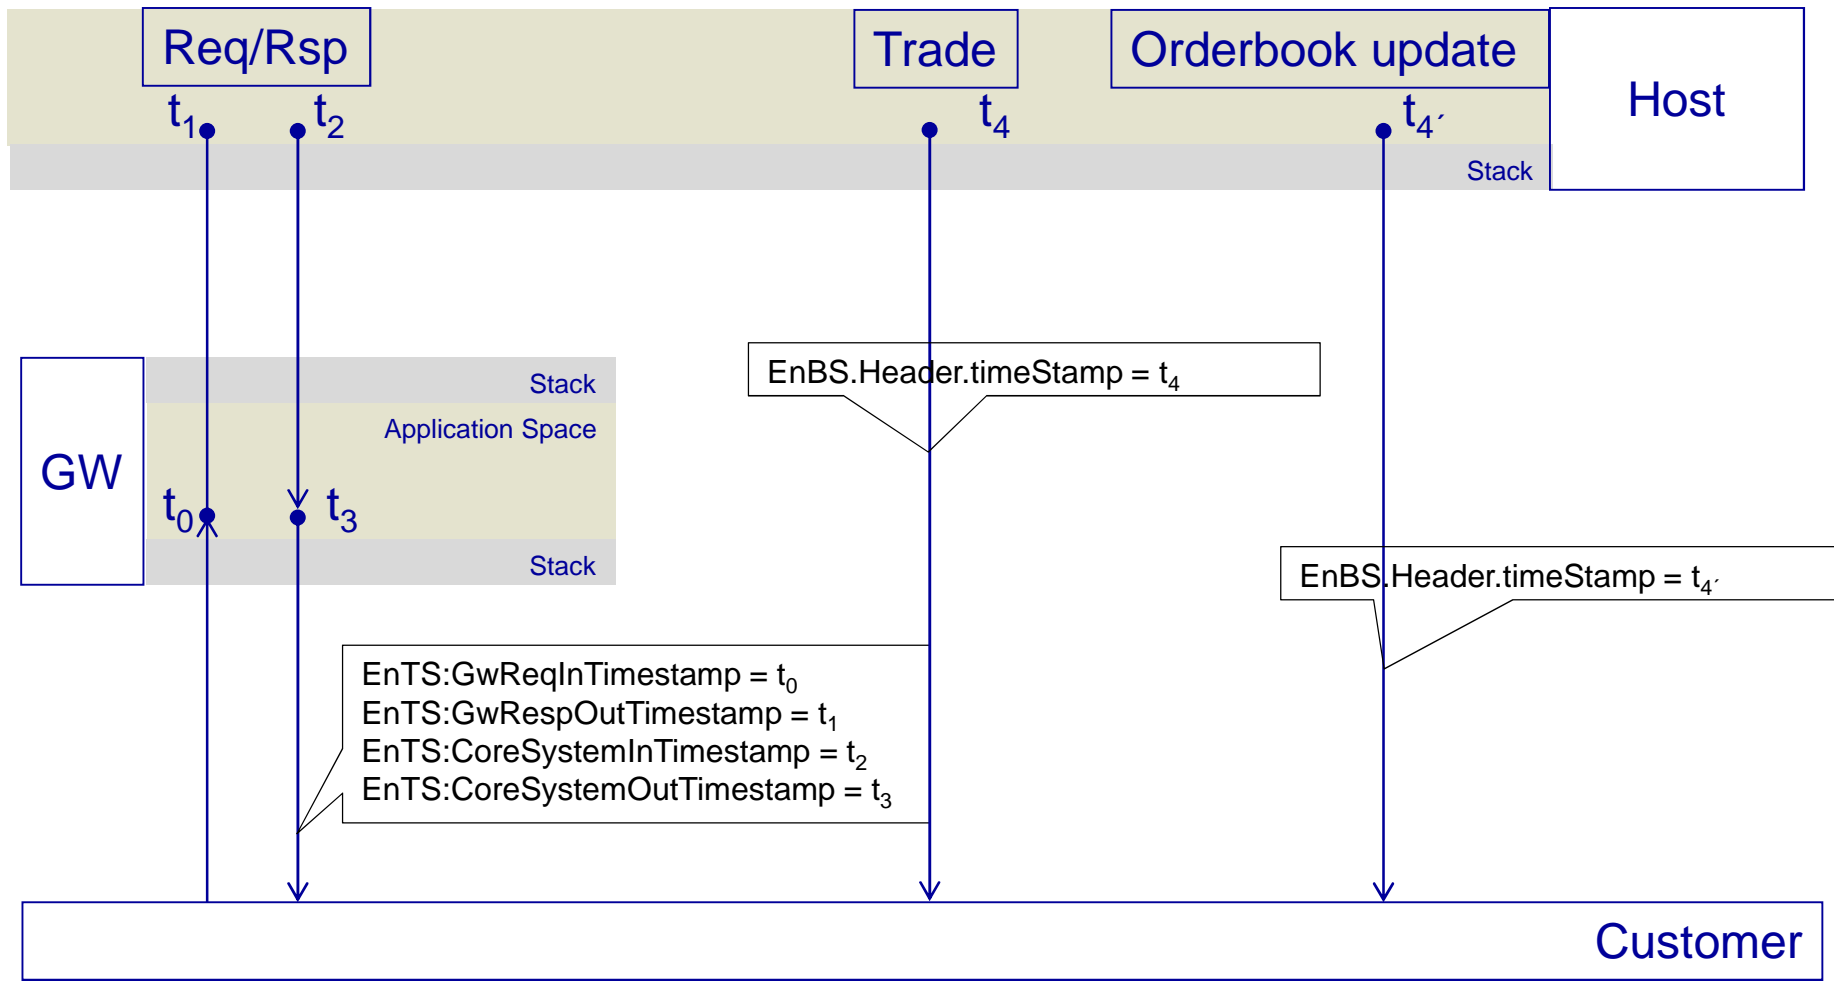
ETS.lastEventTranId is in 100ns since 17 Nov.1858

This is an indication of unconsolidated development in the past.

The purpose of the next few slides is to give a consistent view on the most important timestamps in Eurex Enhanced Transaction Solution and Enhanced Broadcast Solution.

Only “Continuous Trading” will be discussed, no Stop Orders, no Auctions etc.

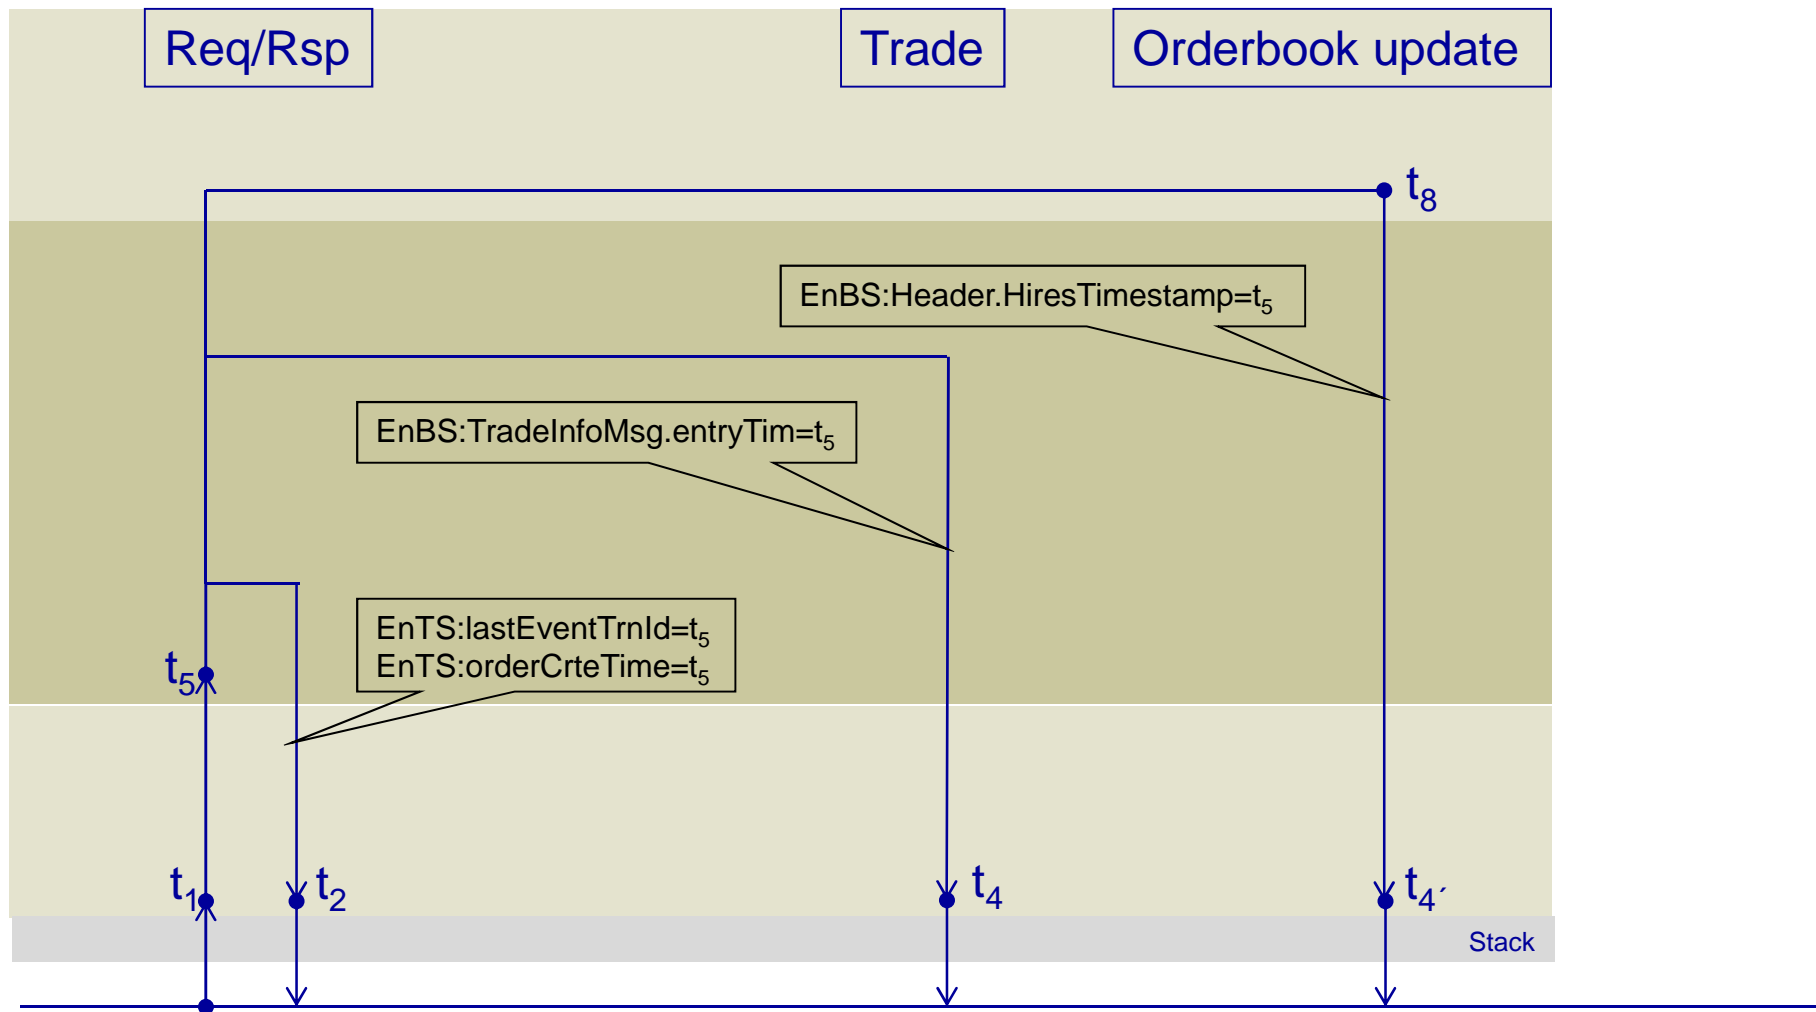
Appendix: Eurex[®] timestamps in Enhanced Transaction and Enhanced Broadcast Solution



Appendix: Eurex[®] timestamps in Enhanced Transaction and Enhanced Broadcast Solution

- § All timestamps mentioned on the previous slide are taken in application space at the very boundary of the respective processes
- § Naming of the timestamps is in sync with Wolfgang Eholzer's presentation of the Eurex Technical Roadmap
- § t_4 is available only in the event of a change to the order book (delta msg)
- § t'_4 is available only in the event of a trade
- § t_4 and t'_4 are available starting with Eurex 14
- § t_0/t_3 and $t_1/t_2/t_4/t_{4'}$ are taken from different clocks which are only loosely coupled

Appendix: Eurex[®] timestamps in Enhanced Transaction and Enhanced Broadcast Solution



Contact

Stephan Hoppe

Deutsche Börse AG
Mergenthalerallee 61
65760 Eschborn
Germany

Phone +49-(0) 69-2 11-1 27 71
Fax +49-(0) 69-2 11-61 27 71
E-mail Stephan.Hoppe@deutsche-boerse.com

Michael Neuerburg

Deutsche Börse AG
Mergenthalerallee 61
65760 Eschborn
Germany

Phone +49-(0) 69-2 11-1 56 41
Fax +49-(0) 69-2 11-61 56 41
E-mail Michael.Neuerburg@deutsche-boerse.com

Thank you for your attention!

