

Load Test Report

Moscow Exchange Trading & Clearing Systems

07 October 2017

Contents

Testing objectives	2
Main results.....	2
The Equity & Bond Market trading and clearing system.....	2
The FX Market trading and clearing system.....	4
The Derivatives Market trading and clearing system.....	5
ASTS Gateways	7
SPECTRA Gateways	8
Latency for transactions	8
Transactional FIX gateways of ASTS.....	9
Speed comparison for different protocols for equity and FX markets (comments for information).....	10
Speed comparison for different protocols for derivatives market (comments for information).....	10
FIX/FAST UDP multicast marketdata of the Equity & Bond Market and FX Market.....	10
FAST UDP multicast marketdata servers of the Derivatives Market	13
Subsystem for real time monitoring of the trading system parameters and market activity	13
Index server, market maker server, MOEX web site	14
Test of clock synchronization over PTP (precision time protocol) at stress load	14
Conclusions.....	15
The Equity & Bond Market, the FX Market	15
Derivatives Market.....	15

Testing objectives

1. To verify the trading and clearing systems operation under conditions of peak loading and an increased number of orders and trades. The trading systems of the following Moscow Exchange's markets were tested:
 - a. The Equity & Bond Market;
 - b. The FX Market;
 - c. The Derivatives Market.
2. To estimate the time of order filling and data delivery from the trading and clearing systems at different load levels and software and hardware configurations.
3. To carry out a public testing of a new FX market system with independent trading and clearing engines.
4. To carry out a public testing of derivatives market FAST protocol with full order log and without data batching.
5. To allow third party software developers and brokers to test their systems and estimate the throughput capacity of communication channels to the exchange venues.

Main results

The Equity & Bond Market trading and clearing system

The testing was performed on a production version of ASTS+ trading and clearing system that was released on 20 March 2017.

The table below shows comparative performance in testing in 2016 and 2017. The 'accepted transactions' term means all the incoming transactions that led either to order registration or to successful cancelation of an order.

	Transactions	Orders	Trades
Values reached (units), 2017	108,550,150	60,733,370	1,735,735
Values reached (units), 2016	113,454,567	65,924,897	1,300,000
Max processing rate for accepted transactions (units per sec), 2017	31,500	16,405	1,356
Max processing rate for accepted transactions (units per sec), 2016	41,822	20,862	1,775


Unlike in 2016, in 2017 the production version of the trading and clearing system was used without any modifications to its configuration. The configuration included the

maximum trade engine speed limitation to avoid market data dissemination delays from the gateway layer.

During the testing of 2016 this limit was not set and the maximum frequencies were measured regardless any market data delays.

Since June 2017 the production equity market configuration acts as a resilient cluster with several engines with hot failover. During the testing this cluster was functioning properly and did not have noticeable impact on system performance.

The graph below shows transaction and order frequencies. Deviation from the peak values between 14:20 and 14:50 was simulated on purpose in order to measure the trading and clearing system metrics at close to the maximum speed levels but before the formation of long incoming transaction queues.



Clients generated 9.3% of the transactions.

The peak frequency at 14:34 was generated by the multiple transactions for cancellation of the already cancelled orders in order to test the system speed for this particular scenario. The maximum possible engine speed was not reached due to limitations of the transaction generator. Monitoring of the engine computing threads during the multiple attempts to cancel already cancelled orders has confirmed that such

transactions are processed approximately 8 times faster than normal cancellations, so this scenario does not have any significant influence on the overall system performance.

The FX Market trading and clearing system

The testing was carried out on the new version of the system with independent trading and clearing engines that is scheduled to go live in H1 2018. This new version is 100% backwards compatible with the current production version but requires separate client connections to be established for trading and clearing data and transactions. Server configuration used during the tests was the same as it was expected to be in production. No changes in hardware are planned.

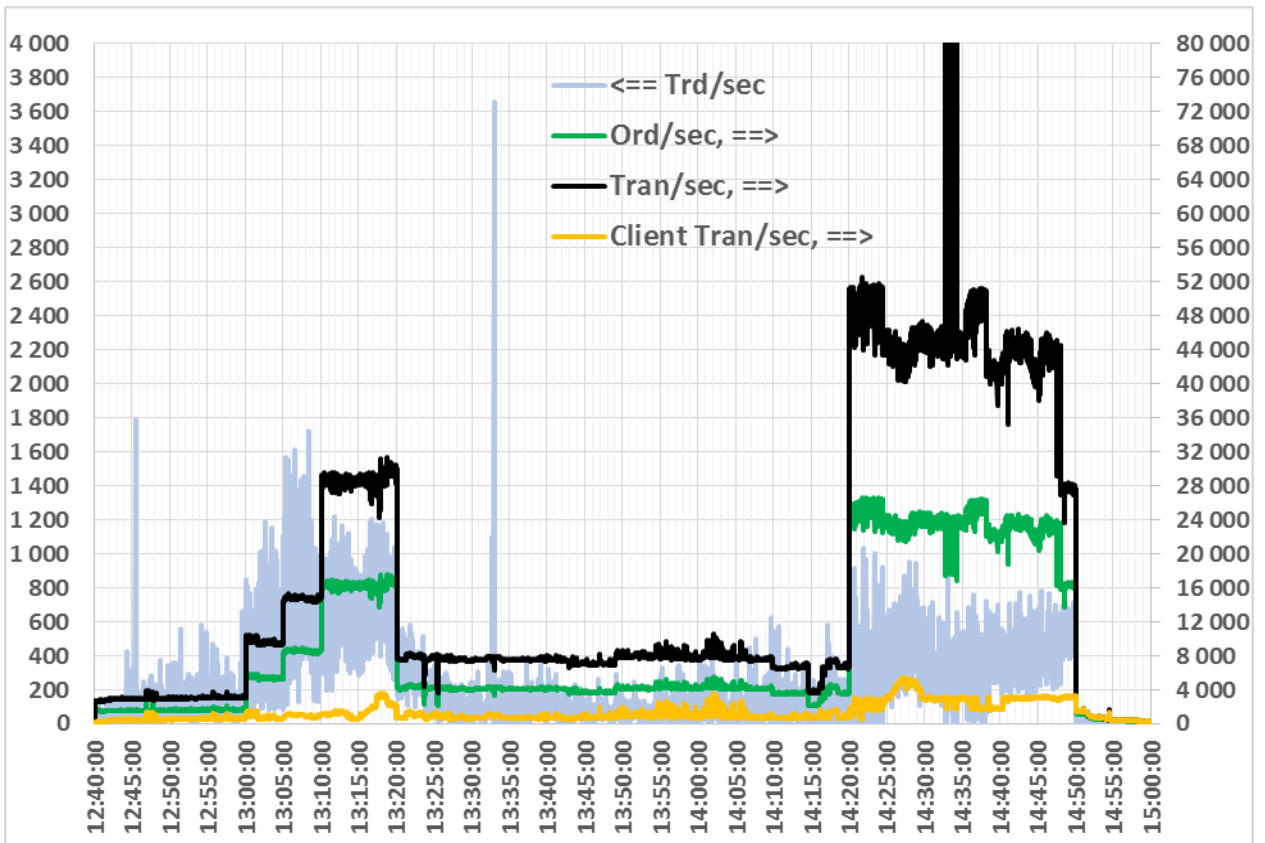
The table below shows comparative performance in testing in 2016 and 2017. The 'accepted transactions' term means all the incoming transactions that led either to order registration or to successful cancelation of an order.

The table below shows comparative performance in testing in 2016 and 2017.

	Transactions	Orders	Trades
Values reached (units), 2017	135,906,705	72,683,460	1,858,811
Values reached (units), 2016	107,447,400	64,061,937	2,583,400
Max processing rate for accepted transactions (units per sec), 2017	52,506	26,584	3,659
Max processing rate for accepted transactions (units per sec), production configuration	30,000	16,000	
Performance improvement, 2017 / 2016, %	+75%	+66%	

There has been a large number of trades generated during the tests – 17 times more than the pick numbers in production during 2017. The maximum system performance levels were not achieved.

The graph below shows the frequency of transactions, orders and transactions by clients – testing participants.



Clients participating in testing generated 8.2% of the transactions.

Architecture of the new version of the system as well as configuration expected to be used in production is free from market data publication delays at top transaction frequencies. Measured performance data is expected to match future production values.

The peak frequency at 14:34 was generated by the multiple transactions for cancellation of the already cancelled orders in order to test the system speed for this particular scenario. The maximum possible matching engine speed was not reached due to limitations of the transaction generator. Monitoring of the engine computing threads during the multiple attempts to cancel already cancelled orders has confirmed that such transactions are processed approximately 7 times faster than normal cancellations, so this scenario does not have any significant influence on the overall system performance.

The Derivatives Market trading and clearing system

The testing was carried out on the SPECTRA system version 5.6 that is used in production since 4 September 2017 on the servers of main DSP data center.


Following the introduction of the TWIME protocol last year the ratio and load profile of transactions submitted over the different protocols was similar to the production.


The order-to-trade ratio in the testing was near the production value. 294 million transactions were send and 12.5 million trades were executed during the testing. The peak performance was 128,000 messages/sec.

	Transactions	Orders	Trades
Values reached (units), 2016	209,830,622	124,200,000	3,500,000
Values reached (units), 2017	294,790,508	159,317,680	12,653,800
Max processing rate (units per sec), 2016	101,000	-	-
Max processing rate (units per sec), 2017	128,000	-	-

During the testing, we carried out the scheduled intraday clearing session. Despite large volumes of orders and trades, clearing was performed as usual within the established time frames.

The graphs below show a transaction load on the Derivatives Market trading system. Clients participating in testing generated 3.62% of the transactions.





ASTS Gateways

Equity market ASTS gateways as well as trading gateways of the FX market installed at both DSP and M1 data centers were functioning correctly.

During the normal operation of the new FX market system the clearing gateways will only be available at the main DSP data center. The clearing engine located at the backup facility works in slave mode and will only serve clearing gateways in case of the main data center failure and migration of the trading engine to the backup facility. This scenario was not included into the load test program.

Clearing gateways were functioning correctly at constant transaction flow rate of up to 45,000 per second. When exceeding this threshold there were delays in clearing data refreshes at gateways as compared to the main clearing engine. In real trading the peak frequencies of more than 45,000 per second happen only in short bursts, so the expected delays will not exceed 10 milliseconds.

ASTS remote Gateways

The ASTS trading platform employs gateways operating in regional technical centers.

To ensure proper operation of gateway servers on each market, at least 8 Mbit/s of network channel bandwidth is required for every gateway software instance of every market.

When constant transaction frequency exceeded 15,000 transactions per second, some remote gateways started delaying data tables if hardware specifications and/or data

channel throughput were not sufficient to sustain increased transaction load. After the end of the project of updating technical facilities in regional data centers that would be fixed.

SPECTRA Gateways

During the load testing no deviations from the normal gateways performance has been witnessed. It should be noted however, that in real trading these gateways serve many more client connections than had been established during the test, so no forecast on server capacity could be made.

Latency for transactions

Equity & Bond Market and FX Market trading systems

Transaction generators at the Exchange side were using Linux version of the embedded ASTS Bridge (libmtesrl.so library) or FIX protocol. These generators had been set up on a server connected to the trading network. Latency data for the FIX protocol for equity and FX markets was collected using the network monitoring system based on Corvil solution and customized for ASTS FIX messages.

Below is the table that shows transaction roundtrip latency for ASTS Bridge and FIX protocols for equity market. Data was collected during the intervals with activity of up to 200 trades per second. For higher trade frequencies median values remain the same but delays for the best 99% increase significantly.

Equity market	Median, Bridge / FIX, microseconds	99%, Bridge / FIX, microseconds
Transaction frequency / time to receive reply		
< 15 000	330 / 312	500 / 470
16 000	350 / 330	600 / 580
26 000	380 / 377	950 / 900
31 000	650 / 650	1500 / 1500

Next is the table that shows transaction roundtrip latency for ASTS Bridge and FIX protocols for FX market. Data was collected during the intervals with activity of up to 200 trades per second. For higher trade frequencies the median values remain the same but delays for the best 99% increase significantly.

FX market	Median, Bridge / FIX, microseconds	99%, Bridge / FIX, microseconds
Transaction frequency / time to receive reply		
< 20 000	300 / 272	400 / 370
30 000	360 / 340	1000 / 870
41 000	410 / 400	1300 / 1250
52 000	700 / 700	1500/1500

As it can be seen the FIX protocol is faster than the Bridge protocol for speeds when there are no long transaction queues at the trade engine. For maximum transaction frequencies latencies align.

Derivatives Market trading system

Internal Spectra monitoring system, Corvil solution and log files for client transactions have been used to measure the derivatives market latency

Derivatives market	Median, TWIME, microseconds
Transaction frequency / time to receive reply	
< 3 000	125
30 000	232
60 000	313

During the testing, door-to-door latency for transactions was less than 500 microseconds on CGate gateways. During the high load with frequencies of over 80 000 orders per second there could be delays in response to order submission over CGate.

For the forecasted peak frequency of 20,000-30,000 transactions/sec next year, the most likely response time (median) will be 230 microseconds for the TWIME protocol with 99% of responses coming not later than in 500 microseconds. At the same time, the response time may reach 2,000-5,000 microsecond in periods of high simultaneous activity of users due to the specially defined restrictions in the trading system.

Transactional FIX gateways of ASTS

This was the first time when the Exchange transaction simulators that work over the FIX protocol were used. When joining these transactions with client activity:

- 24% of transactions on equity market were entered over the FIX protocol. This is close to real trading where 40% of orders are entered using the FIX protocol.
- 20% of transactions on FX market were entered over the FIX protocol. At transaction frequencies of up to 10 000 per second the share of FIX orders exceeded 75%. This is close to real trading where 95% of orders are entered using the FIX protocol.

Equity market FIX gateways setup was the same as in production. They were functioning correctly during all transaction load levels.

For FX market FIX gateways have been reconfigured to work with the new version of the trading system. 9 out of 10 gateways in two data centers were functioning correctly throughout all the tests. One of the gateways in DSP data center failed during the second half of the tests due to error in its configuration that led to insufficient free drive space for log files storage.

Measurement of incoming FIX messages latency was done for 3 equity market and 5 FX market servers located at DSP data center. Corvil equipment was used to measure processing time of incoming FIX messages with types 35=D,F,G on their way to matching engine. For that, incoming FIX transaction timestamps and internal protocol FIX gateway to matching engine transactions timestamps were recorded on the switches serving the FIX gateways. No differences in transaction processing time depending either on particular servers or on types of transaction type have been identified. Also there were no differences in processing time between the client transactions and test transactions generated by the Exchange.

Speed comparison for different protocols for equity and FX markets (comments for information)

No special tests to compare speed of ASTS Bridge vs. FIX protocols have been included into the load test program because of the uncontrolled user activity during this event.

However, a set of such tests have been performed on the production infrastructure in May 2017. They have shown that FIX protocol provides faster delivery of transactions from the colocation facility to the trading engine as compared to ASTS Bridge protocol in 99% of cases.

Using FIX protocol for trading is recommended for all clients who are sensitive to one way and roundtrip transaction latency.

Speed comparison for different protocols for derivatives market (comments for information)

No special tests to compare speed of Plaza2 (CGate) vs. TWIME protocols have been included into the load test program.

As of the Plaza2 architecture the CGate bridge has an internal transport bus (router) that implements authentication, encryption, data receipt and transmission. Since the TWIME protocol does not have such a bus nor encryption, this protocol provides significantly higher speed of data delivery to the Spectra trade engine (30-50 microseconds).

It is recommended for all latency sensitive clients who are concerned both with the transaction delivery speed to the Spectra engine and the minimum delays in replies receipt to use the TWIME protocol for trading.

FIX/FAST UDP multicast marketdata of the Equity & Bond Market and FX Market

Equity market FAST servers configuration was the same as in production. Servers operated correctly during all the load levels.

FX market configuration was customized to work with the new version of the trading system. These servers operated correctly during all the load levels, too.

Statistical data on new order messages from the trading system (OLR feed, 279=0) relative to FIX protocol order accepted messages (Execution Report / 150=0) has been collected using the Corvil equipment.

Since the average order response time is statistically less when orders are submitted over the FIX protocol, statistics on market data publication in FAST feeds relative to replies to FIX orders provides the best metric of the FAST service publishing latency. This data is being constantly collected during the normal trading. There are plans to make this information publicly available.

For the equity market the average time of publication in the FAST feeds as compared to FIX Execution Report during the load levels of 500-25000 transactions per second is as follows:

- FAST, first line (udp.port 16041) – minus 50 microseconds (i.e. earlier than order reply in FIX)
- FAST, second line (udp.port 18041) – plus 110 microseconds (i.e. later than order reply in FIX)

At maximum transaction frequencies the average relative publication times increase to 20 and 600 microseconds, respectively.

The difference between the average publication times between the two FAST lines due to the use of the accelerated data source for the first line that was launched into production on 20 March 2017. The second line does not use this source in order to exclude the necessity for arbitrage between the two lines. In case of the first line failure the second line may be used as failover source with the above stated latency degradation of 160 microseconds.

The test FX market configuration did not include the accelerated data source which is currently on the final development stages. So the data acquired on publication times will only be applicable to the second line servers after the production launch.

For the FX market the average time of publications in the FAST feeds as compared to FIX Execution Report during the load levels of 500-25000 transactions per second is as follows:

- FAST, first line (udp.port 16001) – plus 50 microseconds (i.e. later than order reply in FIX)
- FAST, second line (udp.port 18001) – plus 70 microseconds (i.e. later than order reply in FIX)

FAST publication latency was gradually increasing to 350 and 400 microseconds, respectively, at 48 000 – 52 000 transactions/sec. Difference between the first and

second lines is due to the production configuration of the second line server that is being slightly slower than the fastest possible in the absence of the accelerated data source.

For the reference, the current production average publication times for the FX market FAST are as follows:

- FAST, first line (udp.port 16001) – 0 microseconds (from -7 to +7 microseconds as compared to the FIX order replies for different 5-minutes intervals)
- FAST, second line (udp.port 18001) – plus 137 microseconds (i.e. later than order reply in FIX)

After the production launch of the new FX market trading and clearing systems including the accelerated data source the estimate times will be as follows:

- publication on the first FAST line at the same time as FIX order reply on the average
- for the second line – the same as received during the load test (+70 microseconds).

UDP multicast traffic reached the following values in each copies A and B:

Feed	FX market, Mbps	Equity market, Mbps
Active orders (OLR)	18	14
Market statistics (MSR)	27	20
other feeds	2	2

In 2018 we plan to implement the new version of the system with independent trading and clearing engines on both equity and FX markets. They will have the same estimated maximum performance. Hence when planning the network bandwidth requirements for 2018 we recommend to use the numbers received during the load test for the FX market for the equity market FAST connection as well.

For information: a separate test has been performed to separate the FX market FAST feeds into two groups of traded instruments having approximately the same trading activity in order to improve publication times for mass events. In this configuration the maximum bandwidth required for OLR and MSR feeds are 25 and 36 Mbps respectively. Such a separation makes sense on the FX market only. Additional research will be conducted to verify the usefulness of such changes in production.

Participants connecting to the service via data distribution channels are recommended to carefully select their subscriptions to data feeds and take into account network

bandwidth as the total combined traffic of two FAST lines of the FX Market and Equity& Bond Market in copies A and B may reach 500 Mbps.

In the production environment the short FAST traffic bursts would most probably correspond to network requirements stated above.

Recommendations given at <http://www.moex.com/a1160> are applicable to each FAST line.

FAST UDP multicast marketdata servers of the Derivatives Market

During the test both the production servers with full order logs with no batching as well as production servers that transmit aggregated order books with data split into 10 ms slices similar to Plaza II / CGate services have been used.

The multicast traffic for one data feed (FEED A) for the full order log with no batching has been as follows:

Load, transactions/sec	Mbps (average)
3000	3.0
15000	10.3
60000	50.5
80000	70.0

At 90 0000 transactions/sec the FAST server that transmits full order log has disconnected abnormally. In order to recover the multicast market data with full order log the transaction load had been temporarily decreased to 4000-5000 transactions per second. The cause of the server reconnection has been determined and the appropriate patch will be released soon after its internal and external testing.

No comparison of publication times of FAST FOL vs. TWIME has been made during the load testing because the target configuration of the "FAST with full order log" service is still to be determined in collaboration with trade members.

Subsystem for real time monitoring of the trading system parameters and market activity

Monitoring for the maintenance divisions and technical support. The monitoring facilities operated well and provided data visualization in graphic form. Message signals were produced in accordance with the established criteria; data was collected to the monitoring database without fails. Operation of the monitoring system did not influence the facility performance.

Corvil network monitoring system. During the testing the Corvil (www.corvil.com) equipment and software, which has been adapted to analyze trading systems network

traffic for all markets has been actively used. This monitoring system is being successfully used during the real trading and is being actively developed to provide more functions.

For the purpose of testing additional traffic was registered as compared to the production setup. That led to the stress load of the system during the data capture and significant number of network messages were lost. More attention will be paid to the resource planning during the future tests.

Data on network traffic published in this document has been collected using other network monitoring tools

Despite the stress load on the data capture subsystem the recorded data includes timestamps taken from the single clock with precision of not worse than 1 microsecond. That made the correct analysis of latency possible.

The network timing monitoring system proved to be tolerant to the stress load:

- it did not fail to operate
- 50-75% of events have been detected correctly with no distortion
- statistics were shown in real time
- after the end of the testing the standard web interface has allowed to successfully export tens of gigabytes of the collected data for further analysis.

Monitoring system overload in real trading is impossible.

Index server, market maker server, MOEX web site

All the listed systems have operated correctly, no failures or performance problems have been noted.

Test of clock synchronization over PTP (precision time protocol) at stress load

For the time of load testing an infrastructure for synchronizing system clocks over a high precision PTP protocol has been deployed. This system is ready for production deployment in 2017. During the tests its stability during the unusually high network loads in the Exchange infrastructure has been checked.

Precision checks of clocks on the network devices have shown that the time deviation has been no more than 500 nanoseconds and that is considered to be the excellent result. No failures of synchronization on network devices or servers have been noted.

Clock synchronization accuracy between the FX market matching engine and the first line FAST market data server was measured during the period of constant load between 13:20 and 14:20. For each second, we calculated an average difference between the

times of a transaction registration in trade engine and FAST protocol Sending Time (52). The average publication time deviated from the horizontal line for less than 5 microseconds, which was most likely caused by variations of transactional activity generated by the users. This evaluation also includes the precision of calls to get system time that are used in assigning timestamps in matching engine and in FAST publishing software.

For comparison: the same evaluation has been made for the second line FX market FIX servers that were using the NTP protocol along with the trading engine that used PTP. Between 13:20 and 14:20 the time deviation was smoothly changing from -600 to -1800 microseconds. Such deviation is within the expected NTP time synchronization accuracy levels.

Conclusions

The Equity & Bond Market, the FX Market

1. Performance of the new FX market system with independent trading and clearing engines is 75% higher compared to the current production version.
2. All the components of the new FX market system kept functioning correctly and with acceptable technical metrics on all levels of transaction frequencies. No failures of the central engines or any subcomponents caused by program errors or overload have been registered.
3. Performance of the production version of the equity market system in 2017 is the same as it was for the release-candidate version in 2016. No failures of the system components caused by program errors or overload during this testing have been registered.

Derivatives Market

1. SPECTRA's performance is sufficient to meet demands of participants even at peak loads with respect of order procession and market data distribution.
2. Bandwidth requirements for customers who use Plaza 2 protocol are the same as last year:
 - a. At least 4 Mbps is required for stable operation of client bridges/terminals per each software instance.
 - b. At least 10 Mbps for a bridge in case of using feeds with full order/trade log (FORTS_ORDLOG_REPL/FORTS_DEALS_REPL).
 - c. At least 50 Mbps for a bridge that receives accelerated replication feeds.

We strongly recommend to check not only the available network bandwidth, but the quality of communication channels to the Exchange as well. Channels with high packet losses gravely affect latency and may lead to significant delays in data receipt.

3. It is recommended to the customers sensitive to the fast market data receipt to use the FAST protocol. The technology of the FAST market data distribution without slicing ('pure real time') is ready for production use with the highest loads of up to 80 000 orders per second.
4. The required network bandwidth for clients who wish to use the FAST service to receive ORDERS-LOG on a "pure real-time basis" is minimum 100 Mbit/sec per feed. To receive two feeds, FEED A and FEED B, or data from more than one market, the 1-10 GBit/sec bandwidth is recommended.