

Load Test Report

Moscow Exchange Trading & Clearing Systems

26 September 2015

Contents

Load Test Report.....	1
Moscow Exchange Trading & Clearing Systems.....	1
26 September 2015.....	1
Testing objectives.....	2
Participants.....	2
Main results.....	2
The Equity & Bond Market trading and clearing system.....	2
The FX Market trading and clearing system.....	3
The Derivatives Market trading and clearing system.....	5
Crash Test for the Derivatives Market.....	7
Tests of recovery from warm stand-by servers for trading systems of the FX Market and the Equity and Bond Market.....	8
Sub-system for real time monitoring of the trading system parameters and market activity.....	8
ASTS Gateways.....	8
SPECTRA Gateways.....	9
Communication channels and remote gateways.....	9
Moscow Exchange Technical Centers in regions.....	10
Gateways to the trading system.....	10
Response time for transactions in the Equity & Bond Market's trading system.....	10
Response time for transactions in the Derivatives Market trading system.....	10
Transactional FIX gateways of ASTS.....	11
FIX/FAST UDP multicast marketdata of the Equity & Bond Market and FX Market.....	11
FAST UDP multicast marketdata servers of the Derivatives Market.....	12
Conclusions.....	13
The Equity & Bond Market, the FX Market.....	13
SPECTRA.....	13

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 ASTS+ trading system of the FX market.
4. To allow developers of customer software and brokers to test their systems and estimate the throughput capacity of communication channels before the exchange venues.
5. To emulate a full data re-replication on the Derivatives Market to estimate the time of big data delivery from participants or from the Exchange.

Participants

According to recommendations of the IT Committee of the Moscow Exchange, the list of testing participants will be published on the Moscow Exchange's website after the load testing completes.

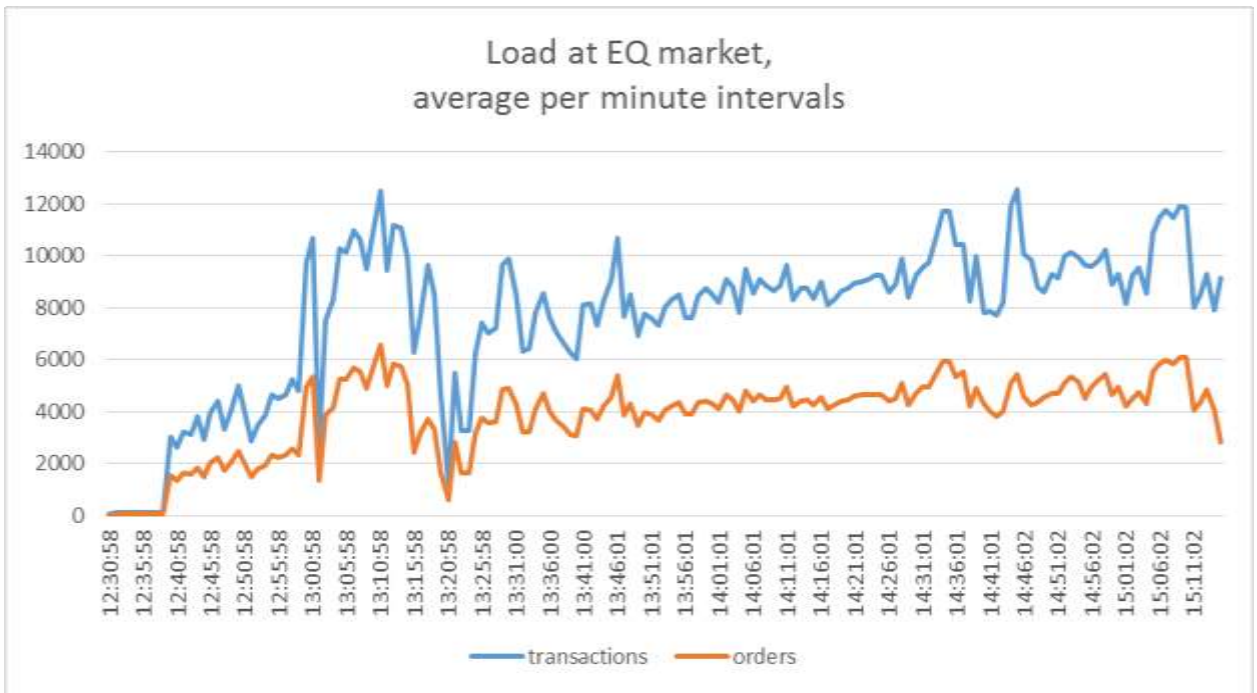
Main results

The Equity & Bond Market trading and clearing system

Performance reached corresponds with expected values for the actual load profile. Testing completed after the trade table was full, i.e. 3 mln trades were made.

	Values reached (units)	Maximum processing speed (units per sec)
Orders	41,500,000	7,826
Trades	3,000,000	770
Transactions	77,100,000	19,314

The graphs below show transaction and order frequency:



Clients generated 38% of the volume. Generating trading activity exactly by clients allowed reaching the more realistic load on all sub-systems of the combined suite “the Exchange-clients” that improved substantially reliability of the results.

During the testing process from 13:15-13:29 and from 13:35-13:41, the trading and clearing system central core was seriously overloaded. Thus, the gateways were overloaded and lagged due to incorrect actions of some participants that entered and remained hundreds of thousands orders at the same price. Such events are not possible in the production environment, however, they disrupted testing for more than one hour.

Participants should apply more carefully the load testing rules.

The FX Market trading and clearing system

The new FX market trading and clearing system that was the system backward compatible in user interface and functionality to the production system, but with much more performance, was used for testing. The total number of transactions, orders and trades were increased to exceed substantially (by 40-50%) values of peak trading activity in 2015 (50 mln order and 75 mln transactions).

The table below shows comparative performance in testing in 2014 and 2015.

The table includes also processing rates for equipment of Data Space data center and the FX Market trading and clearing system for which the generation of excessive data on client positions was excluded as part of the preparation to the upcoming release. We used the trading system code, initial parameters and load generator parameters similar to those used in load testing.

The trading system with positions not considered was not available for using in load testing as customer-side software was not ready for the upcoming release. Processing rates of Data Space are given for reference. However, we need them to estimate appropriately processing rates for 2016.

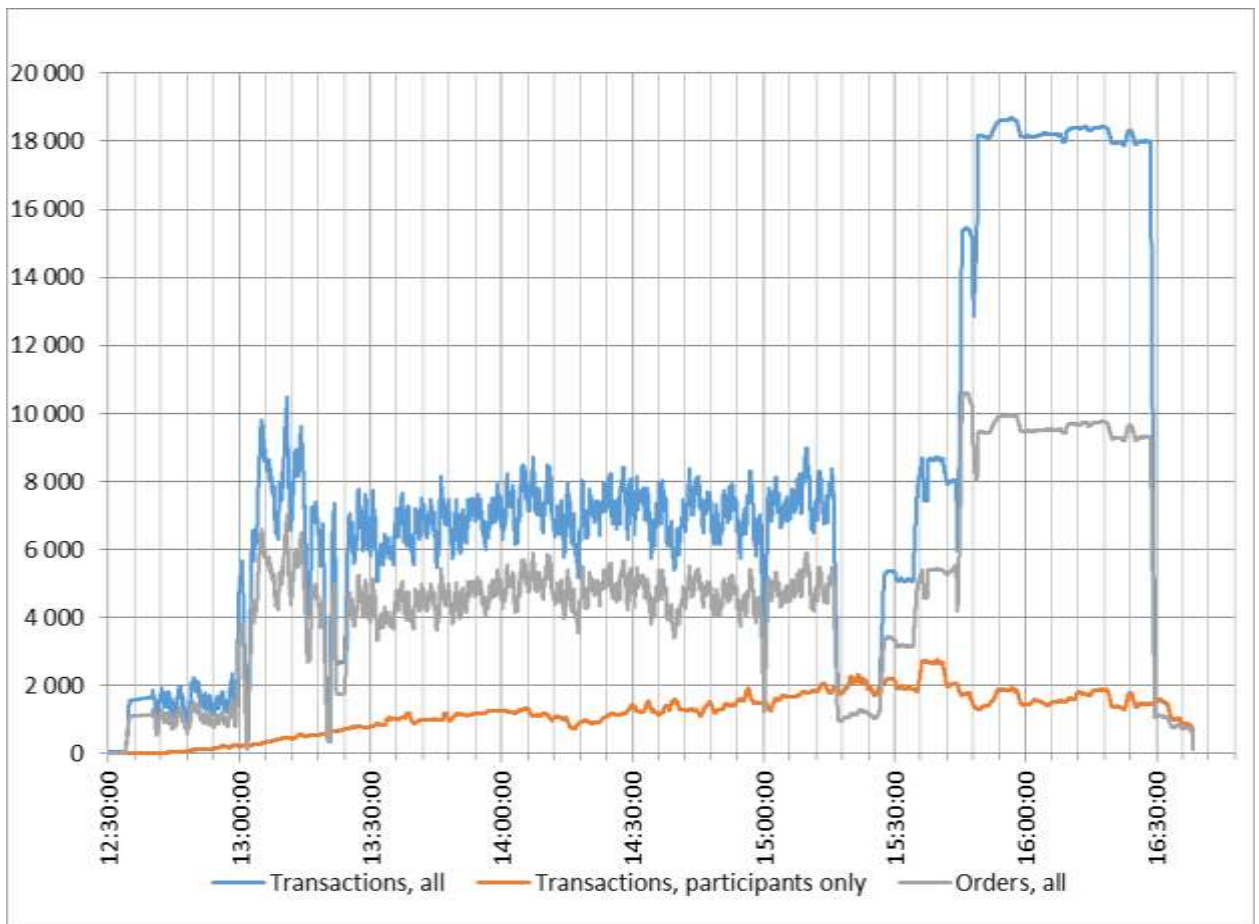
Note: server performance in the trading system central core in Data Space DC is higher by 10-15% than in M1 DC for the same software versions.

	Transactions	Orders	Trades
Values reached (units), 2014	65,568,622	32,994,049	747,539
Values reached (units), 2015	114,790,000	69,994,000	2,898,545
Max processing rate (units per sec), 2014	11,324	5,700	700
Max processing rate (units per sec), 2015	18,884	10,782	1,775
Max processing rate (units per sec), 2015, Data Space, excl. positions for the FX market 3 rd release	35,571	20,100	
Performance improvement, 2015 / 2014, %	67%	89%	
Performance improvement, Data Space 2015 / M1 2014, %	214%	253%	

During load testing, the number of trades exceeded 25x the peak values of the production market in 2015. The peak trade frequency was not reached. Peak frequencies of test trading are given for trades. Such frequencies exceeded 5x-6x peaks of the production FX market. The values in the table are not as great as possible.

During the main phase of testing, we used a load generator with the transaction frequency varying randomly every second from minimum values of the specific average transaction frequency up to 14,500 per second that is approximately corresponded with the production environment. Where the suite operates in such conditions, estimating thresholds and lags of different components of the suite give more accurate results than in the fixed rate thread.

The graphs below show the average frequency of transactions, orders and transactions by clients – testing participants by minutes.



Clients participating in testing generated 15%, at certain moments above 50%, of the total transactions.

The transaction maximum frequency was achieved when only transactions to enter and remove, but not to amend orders were generated. In this case, the comparison of the maximum frequencies of load testings in 2014 and 2015 is accurate. The order maximum frequency was reached when transactions to enter, remove and amend orders were relatively equal that corresponded with the production FX market statistics in H2 2015.

The Derivatives Market trading and clearing system

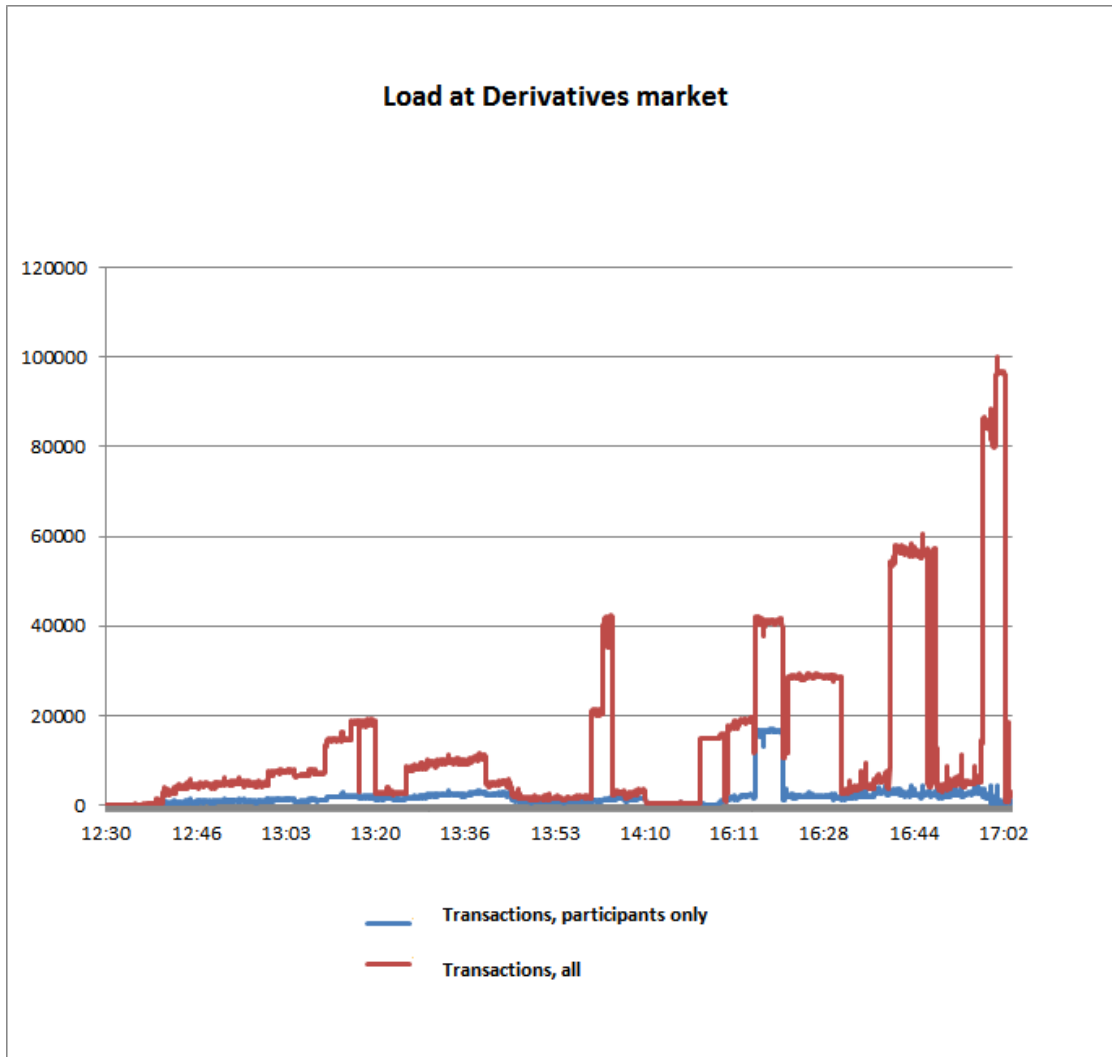
Due to high market volatility and increased turnover in recent months, this load testing of the SPECTRA trading system aimed to imitate a trading day with a load exceeding at least 2x the production load and peak transaction loads added.

Improvements made to the trading system in last two releases allowed increasing substantially the core performance to 96,000 transactions per second without material effect on the RTT latency and market data feed.

During load 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. It is important to note that SPECTRA is not limited in terms of the number of orders and trades per trading day and the figures achieved were determined by testing time and program.

	Values achieved (units)	Max processing rate (units per sec)
Orders	140,000,000	96,000
Trades	19,721,345	21,498

The graph below shows order frequency for the Exchange and participants:



Crash Test for the Derivatives Market

The test aimed to estimate time necessary for a participant to recover in case of an outage on its side or on the side of the Exchange if the outage required data for the trading day to be replicated one more time to the participant's gateways and terminals.

We also tested the ability of the Derivatives Market trading system to switch to the backup infrastructure. During the test, we considered the most difficult scenarios and the necessity to synchronize brokerage systems with the trading system.

Switching to the backup system was successful. It took 60 minutes, and replication of data by gateways took extra 30 minutes.

The test was conducted in conditions most closely imitating peak loads in the production environment. The trading system emulated a hard reset of a life number for main replication streams (orders/trades/order books/positions) when there were 38,000,000 orders and 5,700,000 trades in the system.

The monitoring system and further analysis of the trading members' gateway log files evidences that the data distribution system and participants' infrastructure operated appropriately. The average time of recovery for brokers was nearly 18-22 minutes (recovery in the colocation zone or

in proximity to M1 Data Center took minimum 12-14 minutes) if the Exchange's data distribution servers were connected. Considering the volume of data participants had to receive, the test results were recognized relevant for using to forecast recovery time in case of an outage during trading sessions.

Tests of recovery from warm stand-by servers for trading systems of the FX Market and the Equity and Bond Market

After test trading ended, the system was tested for recovery after an imitated full failure of the trading system main servers.

After the complete halt, the trading system was restarted in the configuration "Main server" + "warm stand-by". After the Main server was restored as at the moment of the trading system halt and we tried to switch the warm stand-by server to the hot stand-by mode, a discrepancy were identified in the number of connection transactions processed. The discrepancy was not critical; it was caused by the automatic switch of the main server to a "suspend" mode after restarting with the warm stand-by not switching automatically to hot stand-by and connections transactions being processed further. Such behaviour of the warm stand-up does not cause any problems and can be eliminated by an explicit call of an operation to transfer the system to the "suspend" mode when the Main server is restarted. The explicit call for procedure to transfer the system to the "suspend" mode was added to the trading system restart scripts.

No problems were identified with the transfer to the "warm stand-by" server when a failure of the Main server was imitated without restarting it preliminary.

The FX market trading system, ASTS+ version, was restarted as usual from warm stand-by, but the start took more time then on current ASTS version. Following the analysis, the ASTS+ system code was changed to improve the start time. Further tests evidenced a substantial reduction of the time of the system restart from log files.

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

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 data base without fail. Operation of the monitoring system did not influence the facility performance.

ASTS Gateways

The infrastructure and gateways at the Exchange's technical centers operated in high capacity providing low latency in updating data tables with frequency reaching 12,000 transactions per second. When market activity increased, updates delayed vs. the trading system central core servers.

During testing, minimum requirements were determined for performance parameters of the gateway computes and communication bandwidth at which delays of gateways were removed. Such minimum requirements will be recorded in a specific document on hardware requirements.

It is important to note that percentage of hardware that did not meet minimum performance requirements was rather high taking into account increased transaction frequency of the FX market trading system. Participants using personal gateways should bring gateway servers parameters in line with the new standards.

Hardware installed at Data Space, the new Moscow Exchange technical center, is balanced in terms of its performance. Some tests evidenced that the trading system facilities and exchange gateways at Data Space operate well at standard rates of data requests and threshold transactions frequency (more than 35,000 per second for the FX market for ASTS+ with partially remoted positions).

SPECTRA Gateways

The infrastructure and gateway servers at the Exchange's technical centers operated as usual without any defects and delays.

To provide environment as close to production as possible, two threads for disseminating market data were risen: standard dissemination (one data bundle per 10 mc) and accelerated dissemination (one data bundle per 0.7 mc). No major faults were identified in the data distribution system. We plan to improve the data dissemination parameters through the FAST protocol optimization of data receipt by users sensible to microsecond delays.

Load testing allowed identifying a potential problem in the use of broker access servers, i.e. all of them failed to receive data at already 39,000 transactions per sec. as any broker gateway server in the infrastructure of the trading system SPECTRA the broker is an intermediate required only to aggregate data and save bandwidth of the channel, the Exchange recommends that participants for whom data speed from the Exchange is important to give up the intermediate in future and increase the channel bandwidth accordingly.

In addition, we elaborated recommendations on improving the backup of SPECTRA at Data Space. We plan to establish two parallel threads to distribute data and accept orders that will be independent to the level of the trading system core. Trading members will be able to connect to both threads simultaneously that excludes connection failures and data omissions in most cases.

Communication channels and remote gateways

ASTS

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

When frequency was more than 10,000 transactions per second, some remote client gateways delayed if computer specifications and/or data channel throughput were not sufficient to sustain increased transaction load. Please see requirements on the Moscow Exchange's.

Participants receiving data via FAST udp multicast marketdata are recommended to decide on the modernization of their infrastructure to exclude also delays of gateway servers due to competition of replication traffic and FAST traffic considering its considerable increase in 2016.

SPECTRA

Due to results of this year's load testing, minimum bandwidth requirements for network remain unchanged:

To ensure proper operation of client gateways, at least 10 Mbit/s bandwidth with the ping time of 1 millisecond maximum is required. To ensure proper operation of client gateways/trading terminals at least 4 Mbit/s bandwidth is required. If a thread includes full order/trade log (FORTS_ORDLOG_REPL/FORTS_DEALS_REPL), a bandwidth of up to 20 Mbit/s and up to 10 Mbit/s is required for access servers and gateways respectively. To get accelerated replication threads, clients should have at least 50 Mbit/s bandwidth outside the trading system.

We strongly recommend that participants check their bandwidth and the quality of their local networks as the network with the high leakage rate deteriorates substantially latency and may cause considerable delays in data transmission.

Please also note that implementation of a backup function in the SPECTRA trading system that is expected in the production environment simultaneously with Data Space, will require double network bandwidth if the function will be used, and ideally, a redundant connection scheme independent at the network level (different providers) will be needed.

Moscow Exchange Technical Centers in regions

Gateway servers at regional technical centers operated as usual when network bandwidth was sufficient for operation without any delays.

Gateways to the trading system

To estimate the time needed for the trading system to respond to orders, we used a generator entering and removing orders via a Linux version of MOEX Bridge (library libmtesrl.so). The generator was run on a server typical for colocation clients of the Exchange that was installed in a trading network segment (the gateway server).

To estimate the time needed for the Derivatives market's trading system to respond, we used a monitoring system, logging of the transaction thread from participants on the side of the Exchange and Cgate API installed in the M1 colocation zone.

Response time for transactions in the Equity & Bond Market's trading system

On the Equity & Bond Market and the FX Market, the average time of receiving a response to a transaction by MOEX Bridge when frequency was up to 15,000 transactions per sec was as follows:

- Mean: 350 microsecond,
- 99% of responses: not later than 840 microseconds.

The average values for the full test cycle are corrupted very much by queues on the trading system's core when ultimate load continued for a long time. They amounted to 750 microsecond.

In tests for ASTS+ at Data Space, the most likely response time (median) was 500 microsecond with 99% of responses coming not later than in 800 microsecond. In this case, we assumed that the peak frequency would be 25,000 transactions per sec in the next year.

Response time for transactions in the Derivatives Market trading system

Door-to-door latency for responses to test transactions was:

- less than 0.2 ms (0.165 ms in average) at load of 39,000 transactions per sec;
- less than 0.3 ms (0.223 ms in average) at load of from 39,000 transactions per sec;
- less than 10 ms (3.8 ms in average) at load of from 80,000 transactions per sec.

The most likely response time (median) was 0.25 ms with 99% of responses coming not later than in 1 ms. In this case, we assumed that the peak frequency would be 20,000 transactions per sec in the next year. At the same time, the response time may reach 2-5 ms in periods of high simultaneous activity of users due to specific restrictions in the trading system.

Transactional FIX gateways of ASTS

FIX gateways operated well throughout the whole transaction frequency range. By the end of testing, fatal delays in trade publishing was seen on two FIX servers operating simultaneously on the Equity & Bond Market and FX Market, due to a shortage of RAM after above 90 mln orders accumulated in two trading systems.

The FIX servers will be divided for separate operation by markets. Considering the proportion of orders on the markets, it is reasonable to install new servers for the FX market and stop using servers with addresses ending with .23 and .24. Data capacity of any replacing FIX servers and a pair of servers with addresses ending with .18 is equal and estimated to be 130 mln orders per day.

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

We successfully tested streamlined configuration of market data servers with the FAST udp multicast technology for the FX market. The configuration performance improved considerably vs. the current production configuration including modes with peak transaction frequencies. When the thread reached 18,800 transactions, threads of anonymous orders and quotes of the FX market FAST first line updated more than 7,000 times per sec with a publication delay less than 200 microsec vs. the response time of the trading system central core to a transaction. The FAST first line used equipment similar to that installed at the new Data Space DC.

Currently, the FAST second line uses less efficient hardware. However, data was published at the same rate as in the production environment. When maximum frequency of 18,800 transactions per sec was reached, updates were published more than 4,000 times per sec with a delay of less than 300 microseconds vs. the response time of the trading system central core to a transaction.

On the Equity & Bond Market, from 13:06 - 14:22, FAST data delayed seriously due to bad initial configuration. After the configuration was changed (the server was not restarted), the FAST first line operated as usual. Publications of the FAST first line delayed at the same rate as the FX Market FAST optimised services until frequency reached 8,000 transactions per sec. When 18,000 transactions per sec were reached values worsened to 300 microsecond with 4,000 updates per sec.

The Equity & Bond Market FAST second line operated continuously, but publications delays increased to 1-2 millisecond.

Speed of FAST services of the FX Market and Equity and Bond Market can be described as follows: when an ordinary limit order is entered, changed or removed, the client in the colocation zone

receives information on such order in FAST threads of anonymous orders and quotes in 100 microseconds in average after the response to the transaction from the trading system.

Improved configuration of FAST servers will be offered in the production environment of the FX Market and the Equity & Bond Market after the final test. We do not expect any other changes in this service.

Some tests at Data Space confirmed that FAST servers prepared for production at 35,000 transactions per sec operated well and data was published with same delays.

Considerable improvement of the ASTS+ performance, new FAST servers and optimization of the FAST service in version 4.0 resulted in higher requirements for the network bandwidth. To set requirements for bandwidth for 2016, it is reasonable to use results of FAST traffic measurement in Data Space with equal relative proportions of transactions to enter, change and remove orders for which maximum FAST update traffic is reached at maximum frequency rate of 30,000 transactions per sec.

In that test, UDP multicast traffic reached the following values in each copies A and B:

- 18 Mbit/sec in every thread for channel of anonymous active orders;
- 18 Mbit/sec in every thread for aggregated quote channels;
- 13 Mbit/sec for market statistics updates;
- 1 Mbit/sec as aggregate for snapshot and instrument descriptions.

Participants connecting to the service via data distribution channels are recommended to schedule carefully their subscriptions for data packages and take into account network bandwidth as the total traffic of two FAST lines of the FX Market and Equity& Bond Market in copies A and B in total may reach 400 Mbit/sec.

In the production environment, peak FAST traffic for the FX market would most probably correspond with network requirements stated above. Requirements for the Equity & Bond Market may be decreased twofold considering its statistics.

Recommendations given above are applied to each FAST line.

FAST UDP multicast marketdata servers of the Derivatives Market

Multicast data distribution servers were tested successfully without any lags occurring even when peak transaction load was reached. Measurements of the speed of receiving trades in the colocation evidenced that receiving data via multicast just as well as receiving trades via native API. We plan to introduce significant improvements to that service such as disabling compression and starting transmitting the full order log.

Traffic in one of the channels A or B was maximum 6 Mbit/sec.

Considering our plans to develop the service and allow backup of data thread by obtaining two separate copies of the data from independent multicast distribution servers, we recommend that participants receiving data via multicast have two separate data transmission channels with at least 20 Mbit/sec throughputs.

Conclusions

The Equity & Bond Market, the FX Market

1. Performance of the FX market's trading system will be increased 1.6x-3x in the upcoming release vs. 2011-2015 parameters. We expect that transaction frequency on the FX market will reach peak values of up to 35,000 transactions per sec or up to 20,000 orders per sec in 2016. We believe also that performance of the Equity & Bond Market's trading system will be improved to the same level in 2016.
2. Gateway servers with performance worse than in models with two processors like Intel Xeon E5 2667 v2 and with memory modules DDR3 1866 do not provide low latency in updating data tables and in responding to requests at peak transaction frequency. The parameters should be included in hardware standards for gateways of the Equity & Bond and FX markets.
3. Network bandwidth requirements for clients using local gateways and/or obtaining market data via FAST udp multicast marketdata must be much higher.
4. It is reasonable for geographically remote clients of the FAST udp multicast marketdata service to consider certain copies of the services with traffic limited to the levels of mid 2015 (maximum 20 Mbit/sec per market in every copy A and B).

SPECTRA

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. We recommend further steps to improve reliability and backup of participants' connections. They will be possible due to our migration to the new data center and infrastructure modernization.
3. Participants are recommended to consider abandoning leased gateways that are an out-of-date part of the "exchange – participant" infrastructure.
4. Network bandwidth requirements for clients using the FAST service must be higher.
5. Gateway servers performance for standard configuration of Intel Xeon series 53XX, 24 GB RAM, Windows server 2008/2012 OS is sufficient to service data requests from clients via Moscow Exchange's gateways when the total transaction frequency is up to 35,000 transactions per sec, if response time remains at less than 5 ms and 30-40 clients are serviced at one time.

New network requirements will be published on the Moscow Exchange's website.