Finteligent Trading Technology Community (FTTC)

Research Report: 10GbE Low Latency Networking Technology Review

Abstract:

Low latency networking is used by electronic trading applications to provide better execution rates and allow more in line processing. The FIX protocol is widely used to implement electronic trades. This paper examines the comparative performance of low latency switches and network interface cards when running the FIX protocol

October 2012

1. Introduction

Finteligent Trading Technology Community (FTTC) is a body focussed on understanding and promoting various technologies which are at the heart of Financial Trading. The periodic research reports which are released by this community are designed to provide a balanced review of the performance of technologies promoted for use in trading. This is the second such report, the first being a review of the performance of FiX engines in the trading environment. An upcoming research report will focus on the performance of CPUs and the impact of deploying trading applications in virtualised environments.

The community attempts to provide 'real world' simulations which can be easily translated into trading environments. FTTC is led by Intel and OnX to provide a forum for discussion and debate and to promote innovation in engineering testing and solution research.

The focus of this report is to differentiate features in the networking infrastructure deployed for high performance trading activities.

2. Summary

The network performance linking computing components of the electronic trading stack has become a critical to the success of firms maintaining a successful 'fill rate' for trade orders placed into financial markets. OnX investigated a range of networking technologies (Network Interface Cards and Switches) which the respective manufacturers claim to have low latency characteristics. This latest research report aimed to investigate whether the deployment of alternative technologies to the components selected in the previous tests FIX Messaging Testing – for Low Latency could impact the latency and performance of the test trading infrastructure. A range of high performance network interface cards (NICs) and switches were tested at Intel's fasterLAB in the London, UK to facilitate testing and comparison.

All 4 of the NICs performed an operation called "Kernel by-pass" which allows an application to write data directly to the network card and skip the operating system. Enabling this process can save between 3-5 microseconds.

The 3 switches had different design patterns. Two used "Cut through" packet processing which forwards the incoming packet to the egress port as soon as the destination address is decoded but before the whole packet has arrived thus enabling lower latency switching, independent of the actual size of each packet. The third switch uses "store and forward" packet processing which loads the inbound packet completely before sending it out, resulting in a variation of forwarding delay as packet size increases. If a malformed or corrupted packet is received an advantage of a store and forward switch architecture is to have the opportunity to prevent it being sent on through the network. Except where errors occur in the frame header, cut-through switches by nature will forward faulty frames, however most implementations mark bad frames to enable efficient dropping by the end station or an intermediate store and forward device.

The test harness we have measured in this paper represents a simulation of a simplified model experienced in trading environments. Matters of scale, matter too. In this respect, we have extended the single switch measurement performed in our previous report to multiple switches, reflecting better how real world systems are deployed. Further work would be required to provide comparative data.

The test harness is an FTTC test suite, developed to generate messages and drive volume through FIX engines to a trading venue. A simulator runs a pseudo exchange/venue that produces between 500 and 100,000 market data messages a second. These data messages are sent over the network to the test rig with price increments of '.001'. When the price ends in '.000' the simple algorithm issues a BUY order. An OMS (Order Management System) function in the harness generates a compliant FIX Protocol message which a commercially available FIX engine exports over the network to the exchange simulator. The simulated exchange performs an instant match and responds with a confirmation message back to the test rig. On confirmation of the BUY, the simple algorithm then creates a SELL order which passes to the exchange simulator for confirmation. Two accurate timings were taken (i) from the market data signal to the BUY order and (ii) the BUY confirmation to the SELL order. A range of market data rates were tested so the engineering team could examine the performance over a range of activity.

The results showed good flat line performance characteristics across the range of activity indicating that consistent network performance was achieved. The latency of the NICs ranged from 4.7 microseconds to 7.3 microseconds. If latency is the only decision criteria there was a clear leader in terms of speed. In order to achieve a rounded perspective on performance jitter and additional features such as PTP (Precision Time Protocol) time stamping and packet capture features were assessed (reflecting the likely required reliability in production trading environments)

The surprising result, in terms of current technology thinking, from the switch test results, was that a store and forward switch from Gnodal performed well (within 100ns) compared to the cut-through switches. The conventional wisdom is that a store and forward switch would not perform with a port to port latency of less than 1,000 nanoseconds. Gnodal achieved an average latency of 483 nanoseconds in our FIX protocol tests.

There is more to the selection of NICs and switches than raw performance and the additional features were investigated. FTTC advises that trading technology optimisation is at heart a lab originated engineering discipline – set against wider commercial and operational considerations.

Acknowledgements

OnX and Intel would like to thank the 5 vendors for their open and collaborative assistance during this exercise, primarily investment in supplying products but additionally engineering input - without this support this report would not have been possible. .

NIC vendors
Emulex, Mellanox and SolarFlare

Switch vendors
Arista, Gnodal and Mellanox

3. Engineering Report

Introduction

As the performance of servers has increased so has the importance of connection between them which has become critical to the optimisation of overall system capability. This level of detail is acutely felt in the context of electronic trading where speed of reaction to the market signals can determine the success of an algorithm's execution of strategy. Time to market, cost, skills and making the right technology choices and trade offs are not easily reconciled and require structured and informed analysis.

This report gives an insight into 2 elements in the state of the art of low latency network computing that can be used for electronic trading.

Given the spread of trading across the market, and increasingly to niche operations on the buy side, not all firms have the staff or time to conduct the detailed technology assessments that OnX and Intel have achieved for this report and both organisations wish to thank the vendors (as FTTC participants) for the willingness to invest and collaborate with the exercise.

The tests were commissioned by David Quarrell, EMEA CTO of OnX and performed at Intel's fasterLAB, London United Kingdom in September, 2012 by Steven Briscoe, Software Application Engineer under the supervision of Evgueny Khartchenko, Senior Application Engineer, both of Intel.

An FTTC test harness asset which generates trades to evaluate the performance of FIX engines was adapted to measure the performance of NICs and switches. FIX is an appropriate workload – being a core element of the trade life cycle and therefore not a synthetic test developed for engineering only purposes. FIX easily translates into the context of the business of Financial Trading.

Many exchanges and electronic trade venues (MTFs) have their own proprietary binary protocols for trading, most also have standards based FIX string based gateways. Market debate focuses on the latency merits of proprietary v FIX de facto standard. Standardisation is seen as a nirvana for the market, and is particularly useful when trading on multiple venues as the time to trading at a new venue is greatly reduced if the firm's trading software does not require special venue specific software to be written and maintained. This paper does not promote FIX over proprietary, but offers useful quantitative data on performance levels that can be achieved. Please read in conjunction with FTTC's FIX engine testing High Performance Trading – FIX Messaging Testing for Low Latency at http://finteligent.net/pg/file/onxenterprise/read/44775

Four NICs were tested

1. Emulex OCe12102-DX 10GbE dual port SFP+ PCI Express 2.0 with FastStack DBL network acceleration. See http://www.emulex.com/products/network-xceleration-nx-solutions/faststack-dbl/overview.html

2. SolarFlare SFN5122F dual port SFP+ PCI Express 2.0 with open onload network acceleration. See http://www.solarflare.com/Content/userfiles/documents/Solarflare_Onload_SFN5122F_10GbE_Adapter_Brief.pdf

3. SolarFlare SFN6122F dual port SFP+ PCI Express 2.0 with open onload network acceleration. See http://www.solarflare.com/Content/userfiles/documents/Solarflare_Onload_SFN6122F_10GbE_Adapter_Brief.pdf

4. Mellanox® Technologies ConnectX®-3 EN (MCX354A-FCB) dual port 10 and 40 GbE with PCI Express 3.0 with VMA™ message acceleration. We ran tests at 10GbE. See http://www.mellanox.com/content/pages.php?pg=products_dyn&product_family=127&menu_section=28

Three (low latency) switches were used. A low latency switch is one that has a port to port latency of less than one microsecond.

1. Arista 7124SX – 1 rack unit 24 port cut-through 10Gb switch. See http://www.aristanetworks.com/en/products/7100series

2a. Gnodal GS7200 – 1 rack unit 72/18 port 10GbE/40GbE store and forward switch.

2b. Dual Gnodal GS7200 Fabric configuration which reflects a highly available deployment. See http://www.gnodal.com/Products/GS-Series/GS7200/

3. Mellanox SX1036 – 1 rack unit 36 40GbE port cut-through or 64 10Gb ports. We ran tests at 10GbE. See http://www.mellanox.com/content/pages.php?pg=ethernet_switch_overview&menu_section=71

4. Method

The test harness software was implemented across the two servers. This hardware rig is resident in the Intel fasterLAB where it is maintained in perpetuity as the baseline for comparative testing of interchanging elements of the stack.

One ran simulators for a market data feed and an execution venue, the other represented a typical software implementation of a real life algorithmic trading system with all applications run on a single server.

The very simple logic of the simulated algorithmic trading component minimises latency and jitter, so allowing the focus of the benchmark to be on the components under test.

No application program tuning was undertaken once tests were commenced. No NIC specific APIs was used, just standard socket APIs. Out of the box command line interfaces with documented CLI switches enabled the use of the kernel bypass drivers.

Governed Benchmarking Process

Each benchmark test is described below, and illustrated in figure

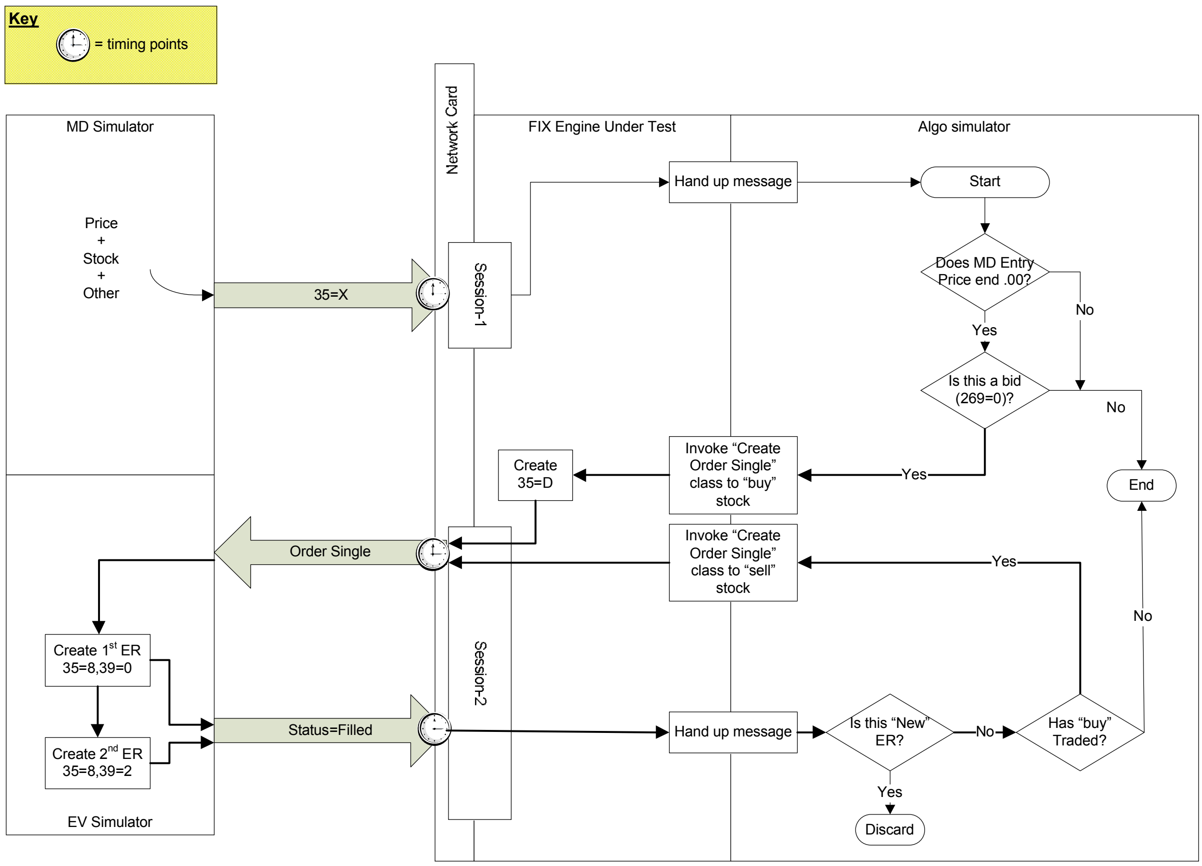

1. The market data simulator created Market Data Incremental Refresh messages (tag 35 = X), assigning an MD Entry Price (tag 270) that was incremented from 0.001 in 0.001 increments, cycling through a small, in memory, list of stocks (tag 55);

2. The FIX engine listened to this stream of messages on a single FIX session (Session-1) and hands each message up to the algorithmic trading simulator;

3. The algo simulator interrogated the data and when a bid (269=0) had a MD Entry Price that ends in ".000" (e.g. 270=56.000) it instructed the FIX engine to create and send a new Order Single (tag 35=D) message to buy 100 lots (tag 38=100) of the symbol (tag 55) to an execution venue simulator on a second FIX session (Session-2).

4. The EV simulator automatically filled the order by creating two Execution Reports (tag 35=8). The first will have an Order Status of "New" (tag 39=0); the second, "Filled" (tag 39=2). These were returned on the same FIX session (Session-2).

5. On receipt of the fill (tag 35=8; tag 39=2) the algo simulator instructed the FIX engine to send another Order Single (tag 35=D) to sell 100 lots (tag 38=100) the same symbol (tag 55).

6. Again, the EV simulator automatically filled the order by creating two Execution Reports (35=8). The first will have an Order Status of "New" (39=0); the second, "Filled" (39=2).

Data Capture

Data was captured by using network taps on the optical fibre network connections. The taps lead to an Endace DAG packet capture network card. After the test run, the packet data is copied from the card to a file on a server for post processing. The latencies of each BUY and SELL transaction were deduced and means and deviations calculated.

The figure below represents the latency of the switch was calculated. We used the same 5m optical cables throughout the testing process. If tests are reproduced elsewhere some variation in the absolute results may occur. However, valid comparisons are made in this paper as each switch and NIC test used identical fibre and latency measurement. The resolution of the packet capture host is rated at 7.5 nanoseconds.

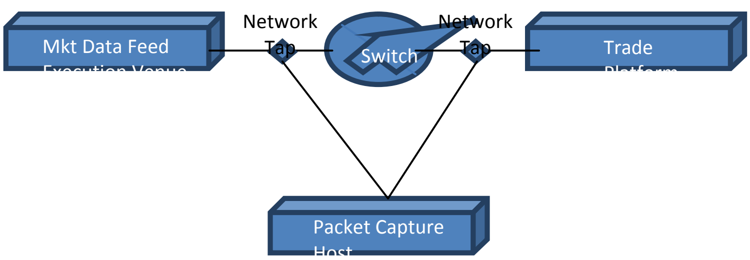

Test rig components

1. Market data and Exchange simulator
   OS Windows 2008 SP1
   CPU 2 x Intel® Xeon® processor
   RAM 32 GB

2. Trade server HP DL380
   OS Red Hat 6.2
   CPU 2x Intel® Xeon® processor E5-2690 @ 2.9GHz
   RAM 192GB

3. Data Capture server
   OS Red Hat 6.2
   Endace DAG 9.2 network packet capture card.
   Note: The Emulex OCe12102-DX with FastStack SNIFFER10G was also leveraged in packet capture tests. Performance numbers covered in Section 6 Discussion, Network Interface Cards.

5. Results

System tuning was an essential step to arriving at the final results with 20-30% improvements between "out of the box" and final runs. Guidance from the manufacturers was useful but proved not to be definitive enhancing input for the specific idiosyncrasies of the electronic trading use case.

Three levels of tuning were undertaken

1. BIOS setting, which turned off all power saving options.
2. Operating System tuning, which turned off unused services, such as FTP.
3. NIC specific command line options, where they were available.

Network Interface card results

The table shows the final results for the NIC tests. Each card was tested with increasing delays in microseconds that showed the performance ranging from 100,000 market data messages per second to 500 market data messages per second. This allowed the engineering team to review the performance over heavy to light work rates – reflecting intraday trading patterns.

Key
BUY - average latency in microseconds from market data signal, a BUY request.
SELL - average latency in microseconds from BUY confirmation to SELL confirmation.
SD – standard deviation.

Range of market data messages from 100k/second to 500/second

	Machine	Network Card	100k	71k	50k	20k	12k	5k	1000	500
BUY	HP Gen 8	Solarflare 5122	6.595	6.625	6.609	6.651	6.685	6.772	7.054	7.056
SELL			6.231	6.268	6.326	6.296	6.301	6.327	6.339	6.43
SD B			0.916	0.896	0.704	0.899	0.768	1.143	0.175	0.16
SD S			0.933	0.887	0.835	0.769	0.772	1.122	1.058	0.207
BUY	HP Gen 8	Solarflare 6122	6.184	6.187	6.232	6.244	6.271	6.3	6.462	6.445
SELL			5.928	5.933	5.949	5.989	6.004	6.031	6.067	6.044
SD B			0.903	0.883	0.796	0.748	0.658	0.569	1.648	0.193
SD S			0.782	0.812	0.678	0.848	0.813	0.973	0.189	0.151
BUY	HP Gen 8	Emulex OCe12102-DX	7.086	7.058	7.087	7.066	7.091	7.1	7.286	7.344
SELL			6.996	6.974	7.116	6.984	7	7.005	7.038	7.082
SD B			1.112	1.044	1.191	0.908	0.762	0.881	0.388	0.381
SD S			2.128	0.827	0.856	0.738	0.726	0.257	0.27	1.296
BUY	HP Gen 8	Mellanox ConnectX-3 EN	4.677	4.663	4.674	4.682	4.739	4.77	4.951	5.003
SELL			4.773	4.745	4.775	4.828	4.881	4.909	4.979	5.007
SD B			1.033	1.035	1.013	0.998	0.948	0.877	0.0495	0.524
SD S			1.007	0.935	1.003	0.854	0.831	1.004	0.497	0.487

The following graph shows the excellent consistency that was achieved by the NICs over the decreasing work load.

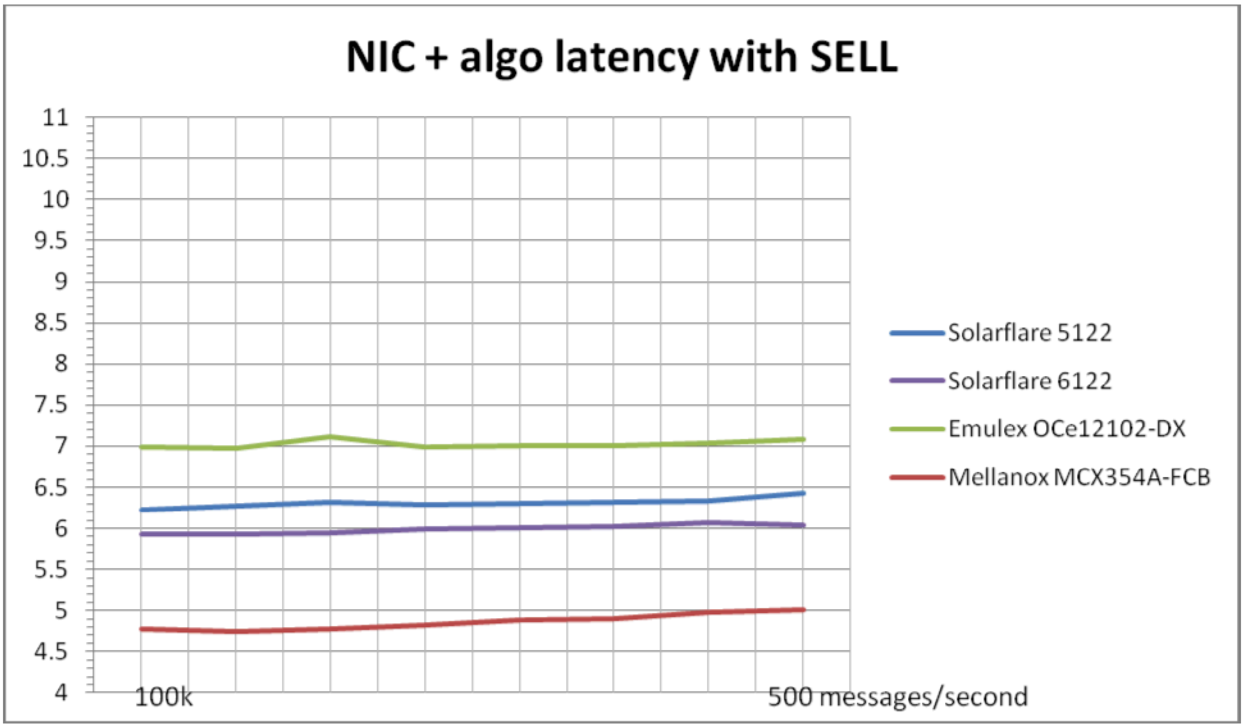

Latency is measured on the vertical axis and the range of messages per second from 500 to 100,000 on the horizontal axis.

The jitter that the different cards induced in the network can be deduced from the standard deviation and is illustrated below where the taller, thinner the spike the better.

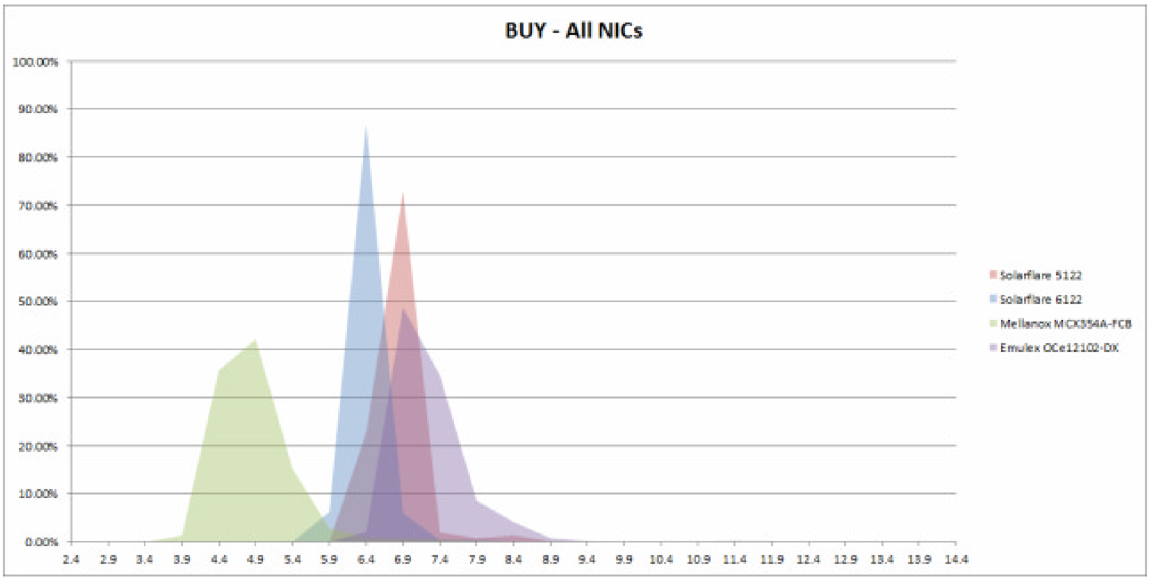

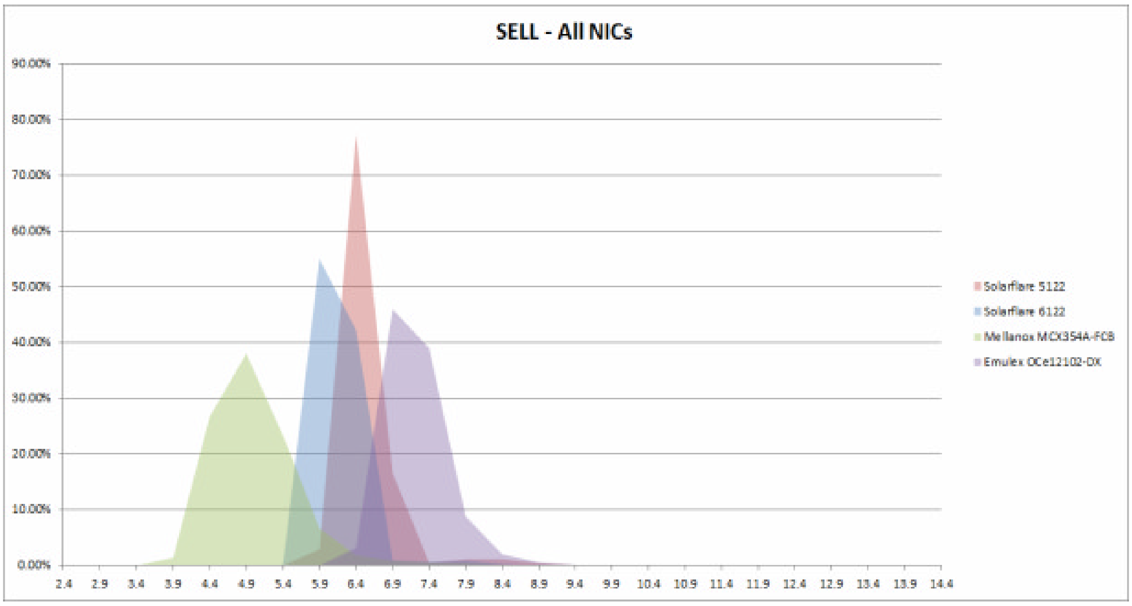

Switch results

The following table shows the final results for the performance tests of the low latency switches.

Range of market data messages from 100k/second to 500/second

Key
BUY - average latency in nanoseconds from market data signal, a BUY request.
SELL - average latency in nanoseconds from BUY confirmation to SELL confirmation.
SD – standard deviation.

Range of market data messages from 100k/second to 500/second

The number of packets measured shows the high number of samples used for calculating the standard deviation and has no other significance.

-	Switch	100k	71k	50k	20k	10k	5k	1000	500
Mean	Arista 7124SX	576.923	576.917	576.894	576.913	576.913	576.952	576.911	576.909
SD		8.912	8.915	8.918	8.913	8.916	8.916	8.914	8.92
Packets		23,091,578	18,769,300	14,970,644	7,233,729	3,885,747	2,016,136	415,594	208,140
Mean	Gnodal GS7200	487.5	487.997	488.343	488.044	487.588	486.724	483.421	483.092
SD		14.878	15.25	15.479	15.659	15.934	16.535	18.307	18.522
Packets		24,271,752	19,411,934	15,216,038	7,299,101	3,891,518	2,019,396	415,936	208,471
Mean	Gnodal over fabric	565.802	563.396	563.287	563.618	563.152	563.036	561.785	561.086
SD		15.398	15.516	15.651	15.433	15.624	15.692	16.478	16.949
Packets		23,349,166	18,878,788	15,026,284	7,247,989	3,885,119	2,018,405	415,458	208,587
Mean	Mellanox SX1036	393.748	394.842	394.013	395.383	396.584	396.653	396.662	396.696
SD		9.379	8.36	9.157	7.723	5.972	5.896	5.844	5.798
Packets		23,482,764	19,159,072	15,010,777	7,267,400	3,894,862	2,018,263	415,642	208,620

The consistent behaviour of the switches of the range of work-loads is illustrated in the graph below.

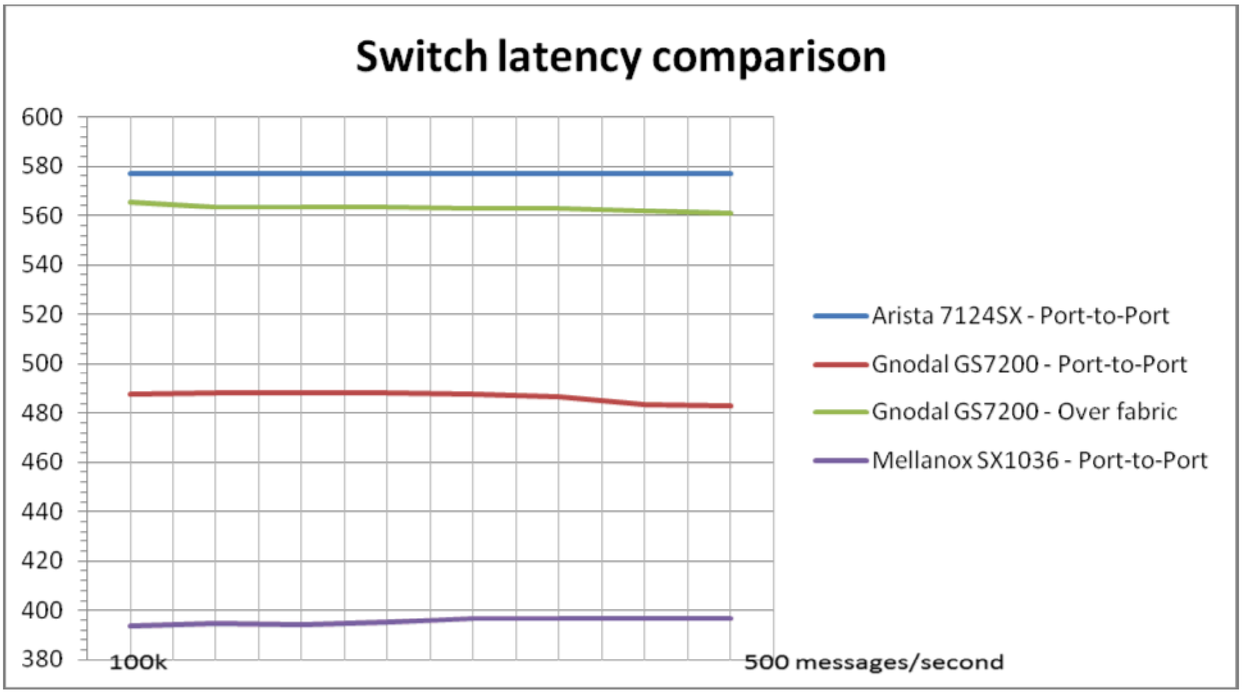

And the jitter of switches illustrated in the frequency distribution graph.


All switches
60%
50%
40%
30%
20%
10%
0%
348 358 368 378 388 398 408 418 428 438 448 458 468 478 488 498 508 518 528 538 548 558 568 578 588 598 608 618 628
Mellanox SX1036
Arista 7124SX
Gnodal GS7200
Gnodal GS7200 (Fabric)

6. Discussion

Network Interface Cards

The investigation of the network cards demonstrated the value of the kernel by-pass which saves 3-5 microseconds of processing time for each packet. The difference between the lowest NIC latency of 4.7 and greatest of 7.3 microseconds can be attributed to differing amounts of processing being done in NIC hardware to firmware.

Additional features of the NICs were assessed as well as the ease of operation. When making a purchasing decision price, ease of use, support cost, as well ease of engagement with the vendor to support current activity and as an indication of likely future innovation are factors guiding the decision. An assessment of each NIC and switch is made below:

Emulex OCe12102-DX ***

The Emulex card proved straight forward to implement and the software acceleration easy to configure. This card was configured and deployed in approximately 4 hours. The latency achieved with this card placed it 4th with a mean latency of 7.086 microseconds. The jitter showed a standard distribution of 1.112 microseconds and compared well to the other NIC's tested.

The ASIC for this card has been in production for 2 years and a replacement is anticipated soon – improvements are expected with the release of their next generation platform.

<u>Additional feature</u>
A unique capability is the FastStack SNIFFER10G software that allows packet capture and injection. This has two applications in the electronic trading environment. SNIFFER10G is available for all OCe12102-DX NICs as a firmware add-on. In this research study, the scalability of the OCe12102-DX was demonstrated by running FastStack DBL and SNIFFER10G concurrently.

The first application of the additional software, which was also tested as part of this research report, allows latency calculation and measurement to be captured without the necessity of an expensive proprietary packet capture HW (like for example Endace card used in the test suite). The measured accuracy of the OCe12102-DX card and SNIFFER10G software was on average 62.06 nanosecond for BUY transactions and 85.17 nanoseconds for SELL transactions. Though it is inferior compared to proprietary solutions (7.5 nanos for Endace) it is sufficiently accurate for the wide range of applications in the trading environment.

The second use which is to capture and replay market data either for post trading analysis or algorithm testing.

Emulex were fast to engage with FTTC and provided the engineering team with a high level of technical support. They have a number of available research papers which make interesting reading and can be accessed at

Mellanox ConnectX-3 EN **

This card was placed first in latency measurement, showing a mean latency of 4.677 microseconds. The jitter is shown by a standard deviation of 1.033. This NIC demonstrated multiple capabilities, including PCI 3.0, 40GbE with a QSA in 10GbE mode (QSA- Quad to Serial small form factor pluggable Adaptor)

The VMA was the most difficult of the 3 to configure but once perfected it was the best for low latency. Configuration and optimisation took the Engineering team over 4 engineering hours to perfect. Of course having learnt the best combination the second and subsequent installations will be a quicker. This was the only card that we tested that runs PCI Express 3.O. The difference between PCI 2.0 and 3.0 is 0.1 microsecond.

<u>Additional feature</u>
This NIC has the ability to switch between InfiniBand and 10/40GbE. It has a wide range of applications within the Financial Trading. A new feature, so new at the time of testing, the firmware was still in beta test, is support for TimeKeeper, which is a commercial PTP client from FSMLabs, and hardware timestamping for all incoming packets. This provides high accuracy PTP synchronization using a standard NIC.

Mellanox is a powerful engineering operation and work at the leading edge of innovation.

SolarFlare SFN6122F **

The latency of this NIC showed a reduction over the previous generation 5122F card. The performance showed latency of 6.184 microseconds. The jitter is shown by a standard deviation of 0.903, the lowest jitter. The network acceleration was the easiest to configure, taking under 2 hours. To have additional PTP capability, that feature has to be implemented with a different card.

Additional feature: PTP capability is available with a different model of NIC from SolarFlare - the SFN6322F. This has the same latency and network features of the SFN6122F that was used for testing, but has the addition of accurate clocks that allow the card to keep accurate time even in the temporary loss of a PTP feed.

SolarFlare SFN5122F ***

The SFN5122F is an older version of the SolarFlare NIC and the card that has been widely deployed in low latency trading infrastructure since 2010.

SolarFlare have to a large extent set the industry standard in the application of NIC technology in trading architecture design – their current products do keep them abreast of newer entrant technology vendors and their road map dictates that they will always be a relevant partner for organisations looking at options within low latency networking deployment.

Low Latency Switches

The headline result from the investigation of the network switches was that a store and forward switch can perform well against the cut-through switches. A cut through switch starts forwarding packets as soon the destination address is decoded even as the packet is still being received on the inbound port. A store and forward switch reads the whole packet before sending out. The advantage is that a store and forward switch will not react too soon and propagate bad packets. A typical store and forward switch will have a port to port latency of 5 microseconds whereas to be in the low latency switch category port to port latency of less than 1 of microsecond is expected. However, current networking technologies provide exceptional inherent reliability and miniscule low bit-error rates (BER), making error handling much less of a headache than it used to be.

Like the NICs, the selection of a switch depends on more than just its port to port latency. Layer 3 routing capability may be crucial as well as the network engineers nemesis, port density per rack unit. Additional features such as analysing port queue lengths and the increasingly relevant issue of PTP capability are also important when deciding which switch to implement.

Arista 7124SX ***

24 10GE ports in 1 rack unit

With its mean latency shown to be 576 nanoseconds combined with support for comprehensive switching and routing the 7124SX rightly holds incumbent position in the electronic trading industry. Arista's mature EOS operating system offers a great blend of enterprise grade functionality and user-extensibility. Arista's engineers added on-board microburst detection with LANZ features to the switch OS to address the needs of trading networks.

While testing Arista announced the 7150S family of switches and the test team were able to get a preview of the product and will report findings later.

Gnodal GS7200***

72 10GbE or 18 40DbE ports in 1 rack unit with 40GbE switch cross connect

The latency associated with this switch is 487.5 nanoseconds. From an engineering stable and relatively new to the trading use case, Gnodal's switch proves that store and forward packet processing and low latency can work well together. The unique fabric capability proved that in multi-switch deployments for both high availability and scalability low latency can be continued. In the fabric test a latency of 563 nanoseconds was recorded between ports on adjacent switches. The lack of layer 3 routing capability confines its deployment opportunity to access layer. As a top of rack layer 2 switch it is hard to beat.

Mellanox SX1036 ****

36 ports 40GbE or up to 64 ports 10GbE in 1 rack unit

The Mellanox switch demonstrated the lowest latency of all of those tested at 394 nanoseconds. Additionally, with 40GbE capability and the ability to run InfiniBand this switch has long legs that will widen its usefulness in Financial Trading environments. There is even the promise that the switch will be able to run both Ethernet and InfiniBand on the same switch. The engineering team particularly liked the easy to use web interface.

7. Conclusions

The switches and NICs tested for this research paper all showed improvement from the results obtained in the previous research paper released in January 2012. This shows that in 10 months, technology released to the market, when harnessed and implemented efficiently, can lower latencies from the average 11 microseconds recorded in January to between 4 and 7 microseconds depending on the configuration selected. Each vendor will undoubtedly continue to release products which reduce the latency of networking equipment further.

The increasing choice of low latency network solutions on the market cannot simply now be made on a narrow 'cost versus latency' calculation. Vendors are differentiating themselves with additional features and functionality which are being engineered into their solutions and offer organisations different options when deploying low latency trading solutions.

When making the decision to deploy low latency networking components, the value of the additional features may have a significant impact on the selection. If 40GbE is going to be implemented in the network within the next two quarters, then Mellanox and Gnodal solutions are attractive. If InfiniBand is in the mix the field narrows to Mellanox. With its comprehensive engineering capability the Mellanox card, Mellanox switch combination has the lowest 10GbE latency and future proofing. In the 10GbE Ethernet trading environments the Arista/SolarFlare combination is hard to argue against as a trusted combination, with many references and relatively straightforward to deploy. The ability of the Emulex adapter to offer competitive latency performance as well as packet capture for latency calculation and market replay/algorithm testing makes the Emulex adapters an excellent choice.

An exciting development and a potential change to the market is provided by Gnodal who have defied conventional wisdom with a low latency store and forward switch which has latency characteristics similar to the cut through switches currently available. Combining 10GbE and 40GbE with its continued low latency capability for high availability and scalability Gnodal is makes a great access layer choice.

The rapid changes in technology and vendor offerings show that this is a fast moving market. This test run was started in mid-September, 2012 and by mid-October both Arista and Cisco announced new low latency switches. The results of this research report suggest that organisations deploying low latency technologies would be advised to partner with a team who can independently test the competing benefits of vendors.

Relatively small increments in latency performance can be gained by deploying technology from different vendors. These may be very important in increasing the execution of a particular trading strategy. Additional features that are provided with these low latency networking components are important to consider to maximise the investment made. Selecting the correct components may minimise further investments to deploy additional expensive infrastructure required.

All participants in these tests collaborate in the FTTC initiatives, and we look forward to providing further insights and advisory capability to the market.