Researched and written by:

Low-Latency.com

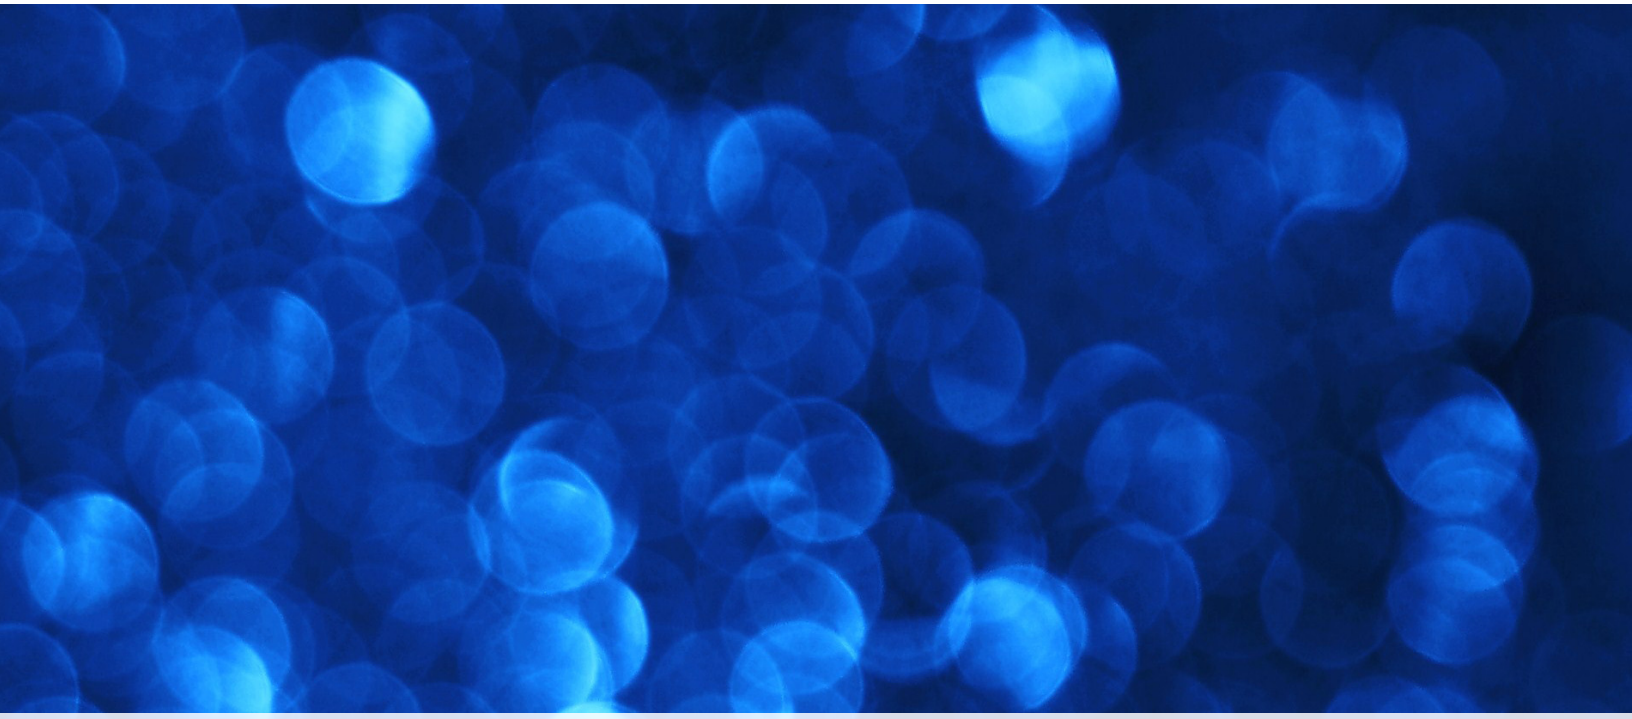

MISSION POSSIBLE:
Turning Data Deluge into Opportunities for Financial Trading

October 2012

An industry briefing prepared by ***Low-Latency.com*** for

DELL
The power to do more

intel

[kx]

KO
KOVE
more from less

The Market Data Deluge

In the financial trading markets, the race to zero is on, with participants including competing marketplaces, algorithmic trading operations, high frequency trading firms, market data aggregators and execution network providers all playing their role in pushing down transaction times from hundreds of milliseconds to tens of microseconds.

Linking all of these players together is data itself, binary digits that have real meaning – and often great value – to those that can capture it, understand it, and leverage it. But as the race heats up, market data rates and volumes are exploding, presenting an extreme challenge to all those needing to manage the deluge, and stay ahead in the race.

It's a vicious cycle of speed. Trading firms seek faster market data to power algorithms, which fire orders faster at execution venues, which are being driven to match those orders more quickly, generating trade reports and impacting order books more rapidly, all resulting in more and faster market data. And this cycle repeats over and over with larger peaks at each cycle high.

Consider this: On October 7, 2011, the consolidated message rate for the major equities, options and futures markets in the U.S. peaked at 6.65 million messages per second (mps), according to marketdatapeaks.com, a website run by the Financial Information Forum.

A large part of the U.S. market data traffic comes from the nine equity options exchanges, which funnel their data to many via the Options Price Reporting Authority (OPRA). In January 2012, OPRA's peak second delivered 4,557,156 messages – a record 49% increase on December 2011 – in a month where trading itself was not that busy.

It all adds up. OPRA sends out several billion messages every day, recording a peak of 14.5 billion in August 2011. It predicts daily message volume to be around 10 billion in 2012, which would require about 1.5 TB of storage if no compression is used.

One might expect message rates to increase when trading picks up, and OPRA is prepared, suggesting it could pump out 7.8 million mps if so required. That figure could be increased to 11.7 million mps in 2013.

With other exchanges around the world looking to increase their matching speed and data capacity, the market data deluge is set to continue for some time. Indeed, statistics for July 2012 show record message rate increases for a number of markets, including NASDAQ (700,613 mps), Bats Options (1,408,545 mps) and Deutsche Boerse (191,058 mps).

Moreover, NASDAQ's recently-introduced FPGA-driven Itch feed - targeted at very latency sensitive trading firms - demonstrates the challenge of managing peak message rates. For this hardware-accelerated feed, the exchange is advising that receiving systems need to be able to cope with bursts of up to 2,000 messages - *per millisecond*.

How Did We Get Here?

Truth be told, market data rates have been on the increase ever since humans were removed from price reporting functions as exchanges automated several decades ago. The advent of electronic systems to support manual trading meant that every quote and trade was captured and distributed as executable data, not just indicative snapshots of the market. Hence data rates and volumes began to rise.

That said, it's been just in the recent few years that data rates and volumes have accelerated to today's extremes. This is illustrated in the figure 1 below, which shows the growth of message traffic – trades and quotes – from just the New York Stock Exchange. Note the exponential increase beginning as recently as 2007.

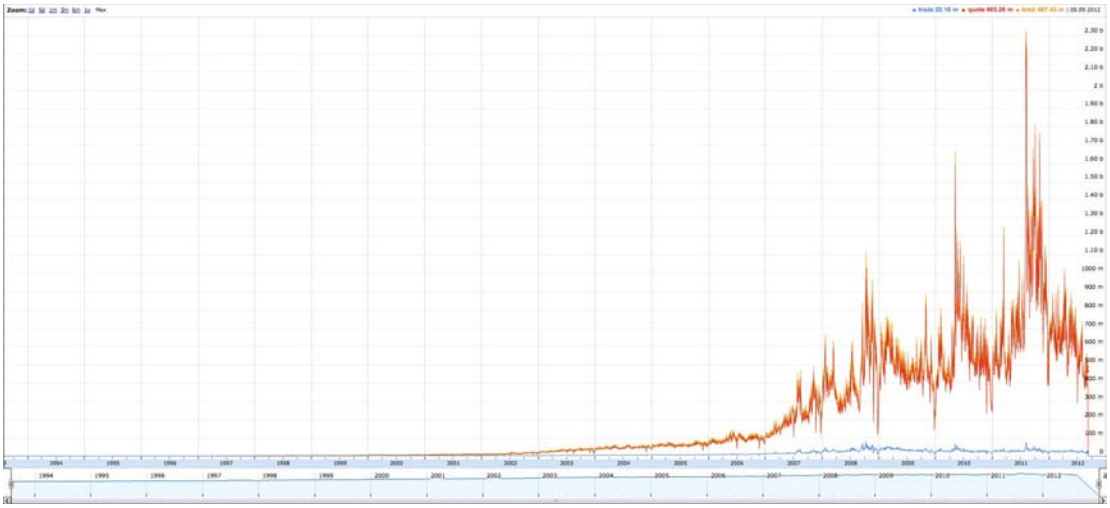

Figure 1: NYSE Trade and Quote Messages 1993 to Present

Some of the key reasons for this are:

- Growth of Algorithmic Trading – which introduced not only automated trading, but in a large part includes algorithms to optimize execution of large blocks, by splitting large orders into many smaller ones, resulting in increased trade executions.
- Advent of High Frequency Trading – by definition involves the very rapid execution of orders to aggregate small profits over time. This generally results in a large number of orders being placed to test the market, resulting in very high message rates.
- Automation of Markets – across all asset classes, from equities, to futures, options, foreign exchange, etc. In the near future, complex derivatives and fixed income markets will see substantial automation. Moreover, this activity is taking place on a global scale.
- Fragmentation of Traditional Markets – as new electronic execution venues have been created – such as Bats Trading and Direct Edge in the U.S. to compete with incumbents NYSE and NASDAQ – which increase competition, especially in the area of transaction times. Faster transaction times results in faster message rates on trade and quote feeds.
- Options Markets – of which there are currently nine in the U.S. with options series of several strike prices each tied to an underlying security. When the price of a security changes significantly, so too do the prices of each strike in the options series. This results in substantial message traffic.
- Technology Advances – especially in areas such as microprocessor speed and multi-core support, and local networking (which underpins co-located market access) – have

allowed markets to execute transactions and distribute data increasingly faster, as demanded by market participants.

- And, last but not least … Volatility – as defined by rate of movement in prices – has been very high in recent years, even though trade volumes have been subdued. This volatility is the product of macro-economic, political and technology trends, which are likely to continue (or increase) in to 2013. See figure 2 below.

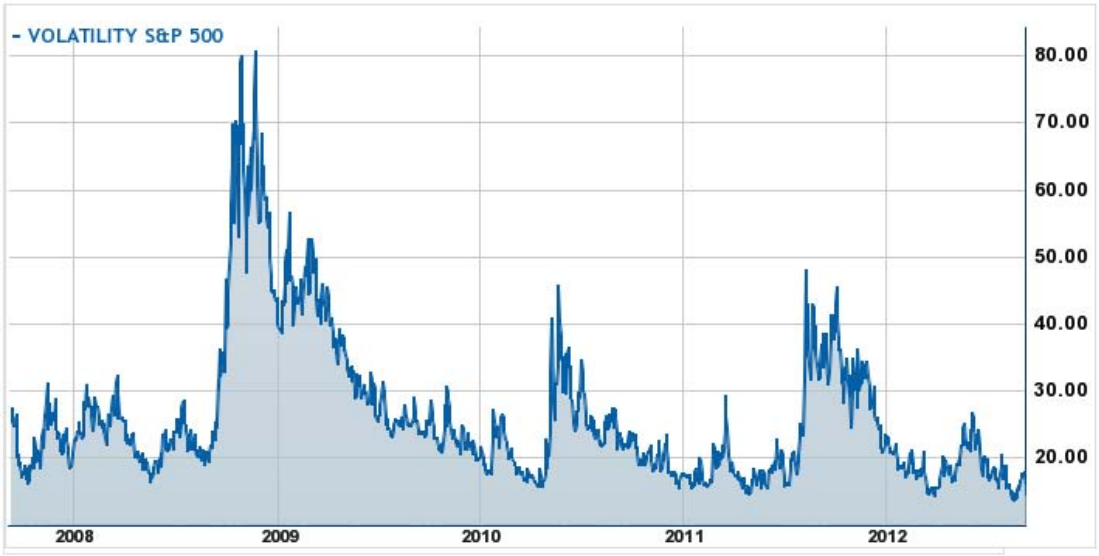

Figure 2: Volatility Year-on-Year (source CBOE)

Thus, the exponential increase in message rates in recent years is the result of a number of inter-connected factors, which have combined to create an on-going 'perfect storm'.

Current indications suggest that the storm is set to continue, even if specific factors – such as High Frequency Trading – have less of an impact. This has resulted in an extreme challenge to trading firms looking to design and deploy infrastructure to handle this deluge of market data.

Market Data Processing Infrastructure

As market data rates and overall volumes have increased, trading firms have explored increasingly complex and costly technology approaches to deal with this reality. Broadly, those firms have sought an infrastructure that can:

- Capture every message received, without any drops, even in extreme – so-called microburst – conditions. The loss of even a single message of tick data can invalidate some trading strategies.
- Store required data fields for each message in a time series for as long as possible. In practice, this means at least a full day, more preferably a couple for input to high frequency trading strategies, and much longer – months to a year – in order to develop quantitative strategies.
- Provide a fast and easy mechanism for applications – several concurrently – to access - the time series.
- Backup the time series and recover fast in the event of server or network failure.

Traditional approaches to this market data processing include the following techniques:

- Where possible, deploying multiple servers and splitting data feeds across them, so that each server processes just a single feed or perhaps a subset – e.g. trades only, quotes only, specific symbol ranges. On the downside, this fragments the time series store and not only makes it more complex to access the entire time series but also makes it impossible to join data across trade and quote tables to find, for example, the current quote of each trade.
- Storing just required data fields in DRAM for fast write processing. The main drawback of this approach being limited capacity of DRAM, often resulting in design trade-offs between what fields to retain, and how to scale capacity naturally.
- Leveraging multi-processor servers, with one processor performing check pointing of DRAM data to hard disk store (or perhaps Flash memory) for backup and for extended storage. This generally requires proprietary code, or licensing of an in-memory database package.
- Using hard disk, or Flash, to consolidate time series for subsequent processing by applications. The obvious issue with this approach is the performance drag of retrieving data from hard disk. And while Flash is much faster than hard disk, it is also of limited capacity.

Essentially, designing a scalable infrastructure for the processing of high frequency and high volume market data is complex, and managing the resulting deployment can be challenging and time consuming.

However, recent innovation in the area of high capacity DRAM storage, coupled with high performance servers and InfiniBand networking, has resulted in a robust architecture that is also straightforward to implement and manage. The High Performance Database Solution from Dell and Kove is a real-life implementation of this architecture, and is now described in detail.

High Performance Database Solution By Dell and Kove

The High Performance Database Solution – comprising Dell purpose-built server technology with the Kove® XPD™ L2 – is a combination of server, InfiniBand networking and high capacity DRAM storage, which has been engineered and pre-configured for optimal performance.

The solution consists of:

- A Dell PowerEdge™ Server, containing multiple Intel® processors. The server is also equipped with multiple terabytes of DRAM.
- Mellanox InfiniBand host adapters and a Mellanox switch.
- A Kove® XPD™ L2 Storage System, which features up to 2TB DRAM per unit. Multiple Kove® XPD™ L2 units can be linked together to provide up to petabytes of storage and multiple terabytes/second of bandwidth while maintaining extremely low and deterministic latency.

Figure 3: Kove® XPD™ L2

Principal performance metrics of the Kove® XPD™ L2 include:

- Bandwidth of 28.5 GB/second.
- 11.7 million input/output operations per second (IOPS).
- 5 microsecond response time.

The Kove® XPD™ L2 storage has achieved sustained and deterministic 5 microsecond random storage read and write to the application. The Kove® XPD™ L2 provides the fastest enterprise storage available without constraints due to usable space, disk utilization, I/O profile, or duration of usage. (See STAC-M3 benchmarks below.)

The 5 microsecond read and write times are achieved for a variable I/O workload over an extended time period, and measured from application space rather than from an analyzer or switch. This is the most rigorous approach, demonstrating performance as an application can use it rather than as a benchmarking number. To put this performance in context, consider the typical performance of hard disk, Flash, and DRAM storage represented in figure 4.

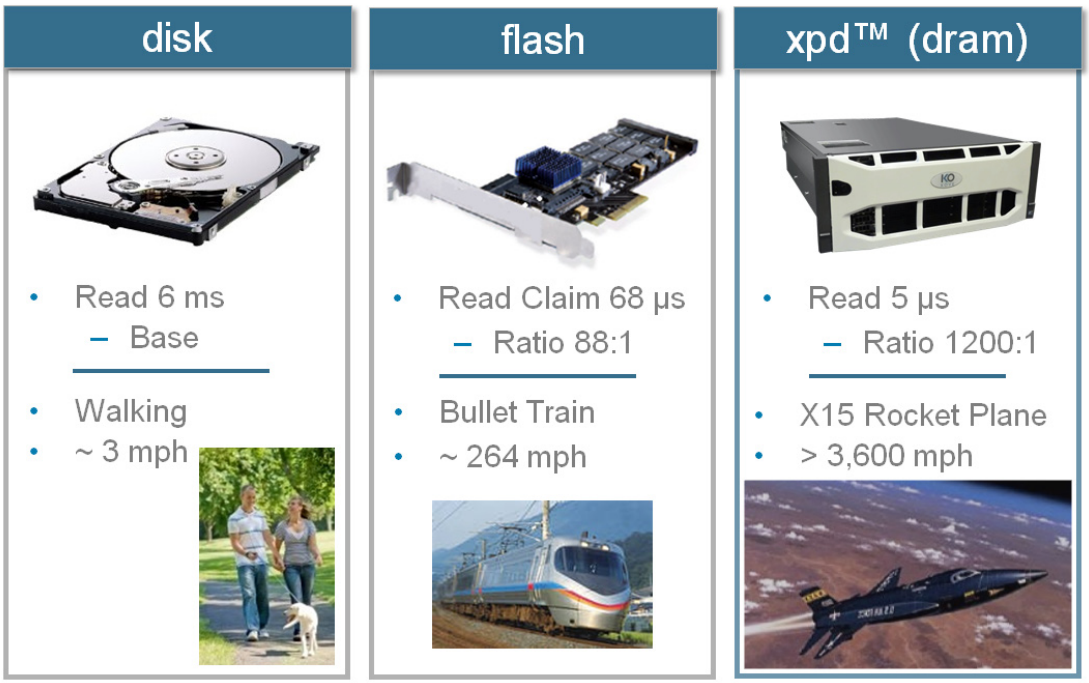

Figure 4: Hard disk/Flash/Kove® XPD™L2 Performance Comparison

Even more important is that this performance is sustainable and consistent for ever – i.e. the performance of the XPD™ L2 never degrades, even as the used capacity increases, over months or years. This consistency is in marked contrast to hard disk and Flash, both of which degrade over time, and as their used capacity increases.

Kove transforms memory into persistent storage, while others try to make Flash storage perform like memory. The XPD™ L2 offers the lowest sustained latency available, not subject to periodic performance jitter or "periodicity". Even under constantly changing disk utilization, the L2 delivers uniform, predictable, and deterministic performance. Such determinism distinguishes Kove® DRAM storage from other approaches, permitting linear scalability into larger, consistent clusters of extreme performance I/O.

The record-setting 5 microsecond latency can remove storage bottlenecks, reduce I/O wait time, and saturate the fastest CPUs. Using conventional network meshes, organizations can scale storage linearly to achieve any performance or capacity level. For example, 4 Kove® XPD™ L2 appliances can deliver more than 46 million I/Os per second (165 billion per hour) and 100 gigabytes per second throughput (360 terabytes per hour). These performance levels leverage DRAM's ability to maintain constant performance levels under load with excellent latency, concurrency, and determinism. Such determinism can be sustained across the full range of I/O, not just a profiled I/O pattern, such as reading 512 byte or 8K block sizes. The memory storage infrastructure provides general purpose infrastructure while still delivering the fastest performance regardless of I/O profile. The Kove® XPD™ L2 delivers correctness of performance across all I/O profiles, across 1 interconnect or many interconnects, across 1 machine or many machines.

The solution architecture used by Dell and Kove allows financial services tasks that used to take hours to complete, to do so in minutes or even seconds. Organizations can accelerate ingest of financial services databases, business analytics, and batch jobs. A great advantage of this approach is that organizations control performance and capacity modularly and linearly, across machine boundaries, simply by adding more hardware components when needed.

The approach's TCO can be tremendous, reducing per-core licensing by using fewer cores to achieve the same results, reducing power and cooling requirements, accelerating pre-existing architectures with faster centrally managed I/O than achievable locally. The solution architecture used by Dell and Kove can be conveniently managed to reduce the cost of delivering exceptional performance. The customer gets to decide what performance is needed and how to service that need, modularly, scalably, and with the industry's least cost footprint. For storage for capacity, spinning disk and Flash may be applicable, but when storage for performance matters, the Dell and Kove approach of using XPD™ L2 storage delivers the industry's best performance and TCO.

Applications needing the ultimate in storage performance can access the Kove® storage directly from user-space, bypassing the kernel for the lowest possible latency, delivering 5 microsecond read / 5 microsecond write times. The Kove® C language API provides a straightforward connect, read, write, disconnect paradigm that makes it easy to integrate into existing or new applications. Reads and writes do not use caching. Applications can perform smaller and larger read/write requests, from one byte to 1 gigabyte. All the while, performance is deterministic and therefore highly scalable.

Other applications – such as for back testing, reference data processing and risk management – might find it simpler to use standard disk I/O functions. In this case, the read time is just 9 microseconds and write time is 11 microseconds, achieving 5.2 million IOPS for 512-byte data blocks (18.7 billion random reads & writes per hour in one box), linearly scalable across *n* boxes.

In a real-world context, a back-testing application saw an increase in performance of 50x simply by replacing hard disk as the time series data source with the Kove® XPD™ L2. No application software changes were required, and no tuning of the Kove® XPD™ L2 or data storage took place. Similarly, a risk computation saw a >20,000% performance improvement. Performance applications regularly receive surprising results.

In addition to DRAM memory, each Kove® XPD™ L2 also includes an array of hot-plug SAS hard drives and an uninterruptable power supply. Together with Kove's Smart Persistence algorithm, they provide data integrity and, in the case of a power outage, gracefully persist all data to disk before a power impact. Essentially, with the reliability as good as or better than hard disk storage, the Dell and Kove approach provides new levels of performance in a new footprint of value. Applications can run faster than ever before possible, or run as fast but less expensively.

STAC Benchmark™ Highlights

STAC-M3™ is a set of benchmark tests defined by a group of financial markets participants, with the purpose of conducting comparable performance testing on time series data stores and the infrastructure that supports them. The performance of the High Performance Database Solution – hosting Kx Systems' kdb+ database – was validated against the full suite of STAC-M3™ Benchmark tests.

In all, some 17 different tests comprise STAC-M3™, mimicking functions and processes commonly performed by financial trading systems, drawing on experience of both the equities and foreign exchange markets. A key metric of all STAC-M3™ tests is the latency of query response, measured for both first result and the final result of a time series query.

The tests vary in focus, from simulating a single threaded query to simulating 100 queries concurrently, on a data set that approximates a full year of NYSE trade and quote data.

Published STAC-M3™ Benchmarks

Documenting the performance benefits applicable to the approach, STAC® audited the STAC-M3™ suite of benchmarks on Dell and Kove systems*. In audited and publicly available tests, the approach used by Dell and Kove with an out-of-the box Kove® XPD™ L2 beat the next best competitive storage system by achieving the fastest benchmarks recorded in 14 of 17 tests**, including:

- over 12x the speed of the next best Market Snapshot benchmark
- over 5x the speed of the next best Theoretical P&L benchmark
- over 4x the speed of the next best VWAB-Day and Week Hi Bid benchmarks
- over 2x the speed of the next best Volume Curves benchmark

These benchmark results illustrate the benefits of applying superior latency, concurrency, IOPS, and bandwidth. Most storage systems cannot provide the necessary concurrency to support high thread load, low latency, and high bandwidth to maximize all phases of this benchmark suite.

Commenting on these benchmarks, Simon Garland, Chief Strategist at Kx Systems, remarked: "I am delighted that we saw major improvements in disk access: all the cores were able to work at maximum speed, rather than having to wait for the drives, as one would in a standard set-up."

The approach used by Dell and Kove with XPD™ L2 storage achieved the fastest benchmarks recorded in 14 of 17 tests. Selecting a few benchmark results can illustrate the benefits of applying superior latency, concurrency, IOPS, and bandwidth.

10 Users, Market Snapshot (STAC-M3.β1.10T.MKTSNAP)

The market snapshot benchmark was designed to demonstrate a trading or analytics system accessing the latest market data for 1% of symbols at a random time. This can be difficult for systems to perform because they must access not only the latest data from milliseconds ago, but also data for low-liquidity instruments that might not have traded for weeks. The large time range stresses systems as they must access recently stored data that may be cached and also data from a random point in the past that must be fetched

from long-term storage. As many symbols in a basket may be simultaneously accessed for correlation analysis and their time range is not known, the latency response is highly significant and can vary significantly. Systems that can perform these correlations quicker will have a competitive advantage in the market.

Analysis of the actual storage I/O during these benchmark runs shows just one spike of relatively light I/O, about 2 GB/s over the course of roughly 2 seconds, showing that this test is extremely latency dependent.

In comparison, the Dell PowerEdge™ server and Kove® XPD™ L2 storage solution beat all previously recorded benchmarks. This benchmark shows the advantage of Kove's excellent low latency and high concurrency, comparable to having your entire storage cached in memory.

10 Users, Theoretical Profit & Loss (STAC-M3.β1.10T.THEOPL)

The theoretical profit & loss benchmark was designed to demonstrate simultaneous random access requests to analyze a basket of trades to determine the theoretical profit and loss. Data for each name in the basket is taken from a time range spanning up to five days of forward access where searching is performed to determine volumes that correspond to 2X, 4X and 20X the trade size. This benchmark benefits from the XPD™ L2 low latency and concurrency.

Analysis of the storage I/O during these benchmark runs shows simultaneous reads across random locations and just one spike of relatively light I/O bandwidth of about 2-2.5 GB/s over the course of roughly 2 seconds. This shows storage access over a very wide spectrum of data, which will significantly exercise virtually any storage device.

In comparison, the Kove® XPD™ L2 and Dell PowerEdge™ solution beat all previously recorded benchmarks. This benchmark shows the advantage of Kove's excellent low latency and high concurrency to deliver consistent performance over diverse I/O profiles, except the L2 is persistent and modularly scalable.

100 Users, Statistics at Unpredictable Intervals (STAC-M3.β1.100T.STATS-UI)

This benchmark was designed to demonstrate the potential of a system to perform sustained, heavy reads at random locations across the database. Basic statistics are created for high-volume symbols for each minute over a 100 minute period that crosses date boundaries where the dates span an entire year. This is typically performed when creating end-of-date statistics for correlated baskets and demands heavy I/O, high concurrency, and large throughput.

Analysis of the storage I/O during this benchmark shows an initial spike (over 12 GB/s) and then medium-to-high I/O levels for most of the test (averaging greater than 4 GB/s over the entire run). The fact that the test has 100 client threads indicates that XPD™ L2 storage is providing excellent concurrency to support that thread volume with very low latency and desirable bandwidth.

The Kove® XPD™ L2 storage and Dell PowerEdge™ solution again delivered the fastest performance for this market trading/analytics benchmark, beating all previously recorded benchmarks. This benchmark also shows the benefit of XPD™ L2's exceptionally low latency and high concurrency, while simultaneously running under high bandwidth load.

10 Users, Aggregate Statistics (STAC-M3.β1.10T.STATS-AGG)

The aggregate statistics benchmark was designed to demonstrate the sustained performance of analytics that compute basic statistics for all the symbols on a single

exchange for a large date range, highlighting the performance of an underlying system that is able to deliver massively parallel throughput.

Analysis of the storage I/O during these benchmark runs shows three phases of I/O activity punctuated by periods of CPU-intensive processing. There is a brief initial spike of activity (6-8 GB/s), followed by a period of steady medium-rate I/O (3 GB/s), followed by a period of high-rate I/O (10-12 GB/s).

Most storage systems cannot provide both the low latency and high bandwidth required simultaneously to optimally satisfy all three phases of this benchmark, which XPD™ L2 storage provides easily.

STAC-M3™ Results Summary

In each case these benchmarks produce demanding I/O workloads. Kove® storage optimally, deterministically, and simultaneously delivers the 3-prongs of high performance I/O: high bandwidth, low latency, and high concurrency. The world's fastest IOPS are also achieved as a byproduct of low latency and high concurrency.

Kove® XPD™ L2 provides the storage industry's undisputed performance leader with Swiss Army Knife convenience, delivering performance that lead's the financial services industry's most rigorous benchmark suite.

Combined with Dell PowerEdge™ servers, the approach used by Dell and Kove surpassed all other approaches in 14 of 17 benchmarks.

NOTE:

* For the detailed STAC Report™, see www.STACresearch.com/node/10175.

** Reflects publicly available STAC-M3™ results as of September 28, 2012.

Integration with kdb+

kdb+ is a combination of a high performance columnar database combined with an analytic query language with functions for handling time series data. It combines in-memory and on-disk storage, and leverages multi-core and multi-socket servers.

With minimal integration, comprising only configuration changes, the Kove® XPD™ L2 can boost kdb+ performance in a number of ways.

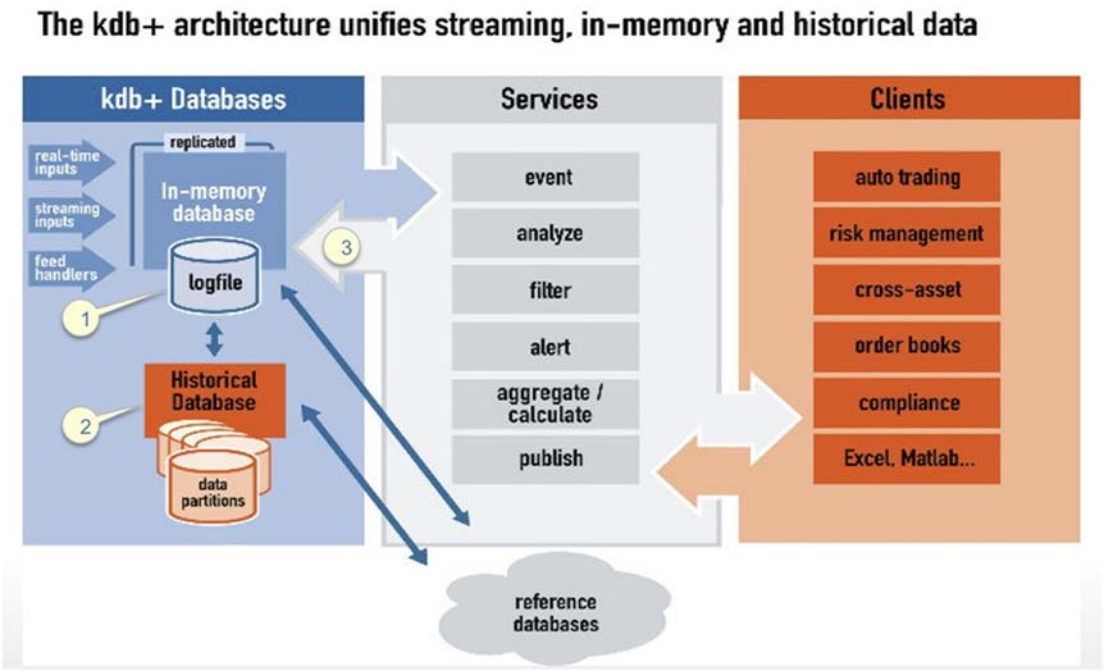

Figure 5: kdb+ and Kove® XPD™ L2

Referring to figure 5 above, the Dell and Kove solution can boost kdb+ performance by:

1. Hosting the log file with tick-by-tick current day price updates on the in-memory database, allowing for extremely fast recovery in the event of system failure.

2. Hosting the historical database of tick-by-tick price data, for use by trading systems operating in the live market.

3. Logging of recent days of tick-by-tick data, as well as trade execution status of trading, risk and compliance systems integrated with kdb+.

Using the High Performance Database Solution by Dell and Kove, users can extend application coverage and make kdb+ run faster in the real world. Applications that previously could not be built can be conveniently built with the new level of performance available to the solution.

Use Cases – Real-Time Market Data and Low-Latency Trading Systems

In-memory and on-disk database systems that support parallel processing and leverage multi-core and multi-socket architectures receive excellent performance benefits from the High Performance Database Solution by Dell and Kove.

- ***Fastest Trade and Order Execution Platform Leveraging In-Memory Performance*** – trading system platforms can be optimized for extreme transaction I/O by leveraging the world's fastest in-memory throughput. An example being the High Performance Database Solution C language I/O interface, which delivers up to 11.7 Million random read and write IOPS in a single 4U unit at deterministic 5 microsecond round-trip latencies, from application to storage and back to application.
- ***Fastest Algorithmic Back Testing of Data*** – the High Performance Database Solution from Dell and Kove supports back testing of large time series (typically 5 to 10 years of tick data) so algorithmic traders can perform these tests in minutes vs. hours, allowing new algorithms to be introduced rapidly, even during a trading day. One global investment bank with terabytes and years of tick data reduced the average algorithm test time from 4 hours to 4 minutes. This allowed for more complete tests that improved the overall trading performance of the algorithm. The ROI for this solution was a matter of weeks, based upon improved algorithm performance.
- ***Real-Time Risk Analytics & Reporting*** – Dodd-Frank legislation and other regulatory requirements, including Rule 15c3-5, have required firms to assess firm wide risk. The High Performance Database Solution provides real-time processing, querying, reporting, and recovery to assess different categories of risk, including credit, market, operational and counter-party. With linear scalability, deterministic I/O, 5 microsecond round-trip latency (application to storage and back to application), and 6.7 petabyte/second instant copies, a new horizon of real-time risk analytics becomes possible. "Real-time" becomes a CPU bottleneck solvable by clustering, and no longer an intractable storage, interconnect, or scaling challenge. For example, addressing out-of-bound trade size or price in the Dodd Frank prevention can be achieved in real-time, scaling linearly when and as needed. Recovery from catastrophic failures can occur in seconds – when time matters real-time.
- ***Real-Time Correlation of Trading and Exposure Across Multiple Asset Classes/ Trading Venues*** – fragmentation of market centers and trading access to both lit and dark liquidity pools across multiple asset classes has become the norm for many securities firms. The resulting challenge has been providing consolidated real time visibility and reporting capabilities that assess overall trading profitability and performance. With linearly scalable memory speed I/O capable of saturating parallelized CPUs, procedures that previously took hours or days can now run in minutes or even seconds at a fraction of the cost.
- ***Real-Time Ingestion and Analytics of High Volume Market/OPRA Data*** – market data volumes are ever increasing and becoming more complex as volume, variety and velocity of data requirements all converge. One increasingly "Big Data" problem is the explosive volumes associated with equities and options data – e.g. a peak of 4.6 million options messages per second from OPRA in January 2012. The High Performance Database Solution can manage these levels of message rates and volume spikes with real time data ingestion, processing, querying and analytics. For example, 4 storage appliances can deterministically deliver more than 46 million I/Os per second (165 billion per hour) and 100 GB per second throughput (360 terabytes per hour) of random read-write I/O. Such design permits compressing ingest landing zones, storage, and query infrastructure

into a single, common infrastructure tier. In turn, this simplifies administration, improves reliability with fewer moving parts, reduces abstraction complexity, and reduces latency to a single tier capable of 5 microsecond latency from ingest on one side to output on the other. Performance, scalability, maintainability, serviceability, reliability, and convenience can all be materially improved.

Further Use Cases – The Possibilities are Endless

Outside of market data handlers and trading systems, the solution can dramatically improve performance for a wide range of applications across a financial services enterprise.

- ***Real-Time Reference Data Access with Concurrency*** – allows for simultaneous access by many applications in pre- and post-trade chain to vital reference data, for improved trade lifecycle processing, leading to lower costs and risks.
- ***Real-Time Transaction Logging and Reporting*** – front and back office activities can be collected and analyzed in real time to provide actionable decision making and reporting. This information can include latency data for different markets, trade fill rates and other trading performance data.
- ***Instantaneous Cloning with Full Back-Up/Restoration of Data*** – the High Performance Database Solution can provide instantaneous cloning and restoration of data. Restoration capabilities have been benchmarked at 1.5 billion rows of data in less than 1 second. This allows for rapid restoration of systems following an outage, minimizing disruption to trading.
- ***Real-Time Predictive Analytics and Reporting*** – the High Performance Database Solution can greatly accelerate performance of predictive analytic software applications, such as Revolutionary Analytics, to greatly enhance the real time collection and processing of information for informed decision making.
- ***Real-Time Analytics vs. Conventional OLAP & OLTP Databases*** – including structured, semi-structured and unstructured data to provide linearly scalable data analysis with significantly higher performance at a fraction of the cost of conventional analytic databases.
- ***I/O Acceleration Front End to any MPP Platform (e.g. ParAccel)*** – the High Performance Database Solution can optimize data environments where massively parallel processing is performed, by ensuring that processing nodes are constantly fed with data.
- ***Real-Time Analytics and Promotion at Retail POS Kiosk*** – leveraging historical customer data, retail banking operations can leverage customer data to create targeted advertising and promotional campaigns to increase banking operation revenues and profits.
- ***Storage Tiering*** – by placing more frequently used data on high performance storage to boost application performance is becoming common practice. But the High Performance Database Solution brings Tier 0 to performance levels that rival in-memory performance, with the simplicity, modularity and scalability of connecting storage devices.
- ***Filesystem Metadata Acceleration*** – can speed up the performance of spinning disk by performing faster metadata access. Examples of filesystems that can benefit are EMC DiskXtender, GPFS, SNFS (StorNext) and HPSS.
- ***Intermediate Application Data Caching and I/O profiling*** – some applications can benefit from fast transient caching of data, but traditional storage can become swamped

by the random nature of read/write operations.

- ***Database Transaction Logging, Table Storage*** – the High Performance Database Solution can greatly improve performance of various types of data bases for mundane yet highly relevant transaction logging and table storage.
- ***Indexing of Big Data Stores*** – placing the search index for vast amounts of data can provide instant search and access.
- ***Acceleration of Mainframe Applications*** – traditional mainframe applications can experience significant performance improvement by back ending, front ending or offloading of specific data processing requirements.
- ***Real-Time Analytics in the Cloud*** – provides acceleration of cloud-based applications, with concurrency providing deterministic performance in multi-tenant environments.
- ***Database Partitioning***– allow administrators to map partitions within databases to different storage tiers using a "partition key", putting hot data on XPD™ L2 appliances and less active data on Tier 1 or less expensive Tier 2 storage.
- ***Database Smart Cache*** – increase database performance by transparently extending the database buffer cache with near-memory speed I/O.
- ***Real Time Anti-Virus and BootStorm Performance*** – in Virtual Desktop environments – organizations running virtualized desktops can avoid delays that come with hundreds or thousands of concurrent users.
- ***Challenging random I/O profiles like strided read/write*** – multiple concurrent data access can create I/O "storms" which can swamp traditional storage, including Flash. The XPD™ L2 appliance has the best "concurrency" properties (ability to handle multiple I/Os at the same time).

Wrap Up and Take Aways

The market data deluge is a reality today with a certainty that it will only increase. Coping with this deluge is an enormous challenge to trading firms, who actually want to move beyond just storing data, but seek to leverage it for their trading advantage.

Traditional systems built to handle high frequency market data are complex to design, require continual upgrades, lack scalability and are difficult and costly to manage.

The High Performance Database Solution provides a high performance platform based on a simple design. In real-life tests it has boosted performance of databases such as Kx Systems' kdb+ by at much as 50x by simply replacing hard disk storage.

The performance of the Dell/Kove/kdb+ combination has been benchmarked and publicly validated by STAC, demonstrating its leading performance in many market data processing functions.

The High Performance Database Solution not only provides an approach to leveraging high frequency market data, but can also boost the performance of virtually any trading, investment or risk management application where data management is a key activity.

Further independent information on market data rates and on the High Performance Database Solution are available at:

www.fif.com
www.marketdatapeaks.com
www.stacresearch.com/kove

www.dell.com

Dell Inc. (NASDAQ: DELL) listens to customers and delivers innovative technology and services that give them the power to do more. For more information, visit www.dell.com.

www.low-latency.com

Low-Latency.com is the premier online community for financial markets technologists who operate at the cutting edge. Launched in 2006, ***Low-Latency.com*** is the portal offering vital news, expert analysis, in-depth research and thought leadership focused on the role played by low-latency technologies in today's high performance electronic trading markets.