



Financial markets participants are rushing to deploy low latency architectures to handle soaring data volumes and to power best execution strategies and algorithmic trading applications.

Designed as a resource for IT managers and systems architects at market participants, A-Team Research's report "**Faster Than A Speeding Bullet – Low Latency Architectures and Building Blocks For Tomorrow's Trading Applications**" explores the market dynamics that have driven the need for low latency systems and how a set of technologies are being deployed to support different trading applications.

Report Contents

- **Market Landscape:**
 - What sparked the interest in low latency – market dynamics and development.
 - Examples of applications requiring a low latency infrastructure.
- **Technology Landscape**
 - The building blocks that make up a low latency market data infrastructure.
 - Architecture diagrams of different technology approaches/applications.
 - Overview of building block/component types and functions with a discussion of performance statistics and how to assess and compare them:
 - o Market Datafeed Handlers
 - o Complex Event Processing and Realtime Databases
 - o Messaging Middleware
 - o Enterprise Data Fabrics
 - Convergence of functional characteristics of products.
 - Vendor Profiles of building block technology providers.
- **Market and Technology Predictions from A-Team.**

Web Portal

To complement the report (and other planned research), A-Team is launching a web portal – www.low-latency.com – as an industry-wide resource for companies seeking advice and guidance in implementing low latency architectures.

The portal will host A-Team's exclusive research, regularly updated news, contributed white papers and articles, podcasts, webcasts as well as interactive elements, including online polls, surveys, discussion forums and real-time moderated chat rooms.

By purchasing this report you can:

- See how Low Latency architectures work
- Consider how to assess and compare Low Latency infrastructures
- Understand how Low Latency can be used to power best execution strategies and algorithmic trading applications
- Learn how Low Latency architectures can help you to manage soaring data volumes


Price: \$495 or £260.

ORDER YOUR COPY NOW

Call Ron Wilbraham at +44 (0)1522 850913 or email ron@a-teamgroup.com

Or visit www.low-latency.com

SPONSORED BY:



**Faster Than A Speeding Bullet ...
Low Latency Architectures and Building Blocks
For Tomorrow's Trading Applications**

WRITTEN & PRODUCED BY:



www.a-teamgroup.com
www.low-latency.com

TABLE OF CONTENTS

Introduction and Foreword.....	5
I. Foreword from Intel.....	6
II. Introduction from A-Team Group.....	8
SECTION 1: Market Landscape and Technology Response	10
SECTION 2: Low Latency Architecture Technologies	19
I. Technology Components and Drill Down	20
II. Measuring Latency.....	32
III. Before You Buy: Questions to Ask.....	36
SECTION 3: Vendor Profiles	38
I. Datafeed Handlers	39
i. ACTIV Financial.....	40
ii. Exegy.....	42
iii. Infodyne Corporation	44
iv. Quant House.....	47
v. Reuters.....	49
vi. Wombat Financial Software	52
II. Messaging Middleware	55
i. 29West.....	56
ii. Reuters.....	58
iii. TIBCO Software.....	61
III. Data Fabrics.....	63
i. GemStone Systems.....	64
ii. GigaSpaces.....	66
iii. Tangosol.....	70
IV. Event Stream Processing Engines	73
i. Aleri Labs.....	74
ii. Cicada.....	77

iii.	Coral8.....	80
iv.	Data Stream Analysis Ltd.....	84
v.	Kx Systems.....	86
vi.	Oracle Corporation	89
vii.	Progress Software Corporation	92
viii.	Skyler Technology.....	97
ix.	SpryWare.....	101
x.	StreamBase Systems	103
xi.	Vhayu Technologies	106
xii.	Xenomorph.....	108
SECTION 5: Appendices		111
Appendix A: Proximity Services		112
Appendix B: About A-Team Group		114

Introduction and Foreword

I. Foreword from Intel

How to Stay Ahead in the Low Latency Arms Race

The low latency "arms race" has been gathering momentum for the past couple of years and shows no sign of slowing down. For traders engaged in certain strategies – such as statistical arbitrage – it is an all or nothing scenario. If they have the end to end infrastructure with the lowest latency, they will spot the opportunities first and trade before their competitors. If they haven't, they will be traded against! Therefore the search is on for the latest technologies that can provide an edge in executing these trading strategies: cost almost becomes a secondary consideration.

But the front office's relentless experimentation with and investment in low latency infrastructure in pursuit of competitive edge is not without its issues. As they continue to push the envelope to achieve ever greater levels of speed, are firms opening themselves up to new operational risks, for example? And is it time to start applying cost benefit analysis to low latency investments, to ensure the high spends are genuinely justified by business benefits?

There has also been a tendency for institutions to address low latency in a piecemeal fashion, focusing on individual components of the low latency stack rather than optimising the stack in its entirety. There is little point in putting "go faster" stripes on a car if the engine isn't tuned properly. If, as expected as a result of the Markets in Financial Instruments Directive (MiFID) in Europe and Regulation NMS in the US there is increased liquidity fragmentation, and firms start to search for best price in real time, this will be an additional driver to take a more sustained approach to achieving low latency than we have seen to date. Key requirements for the next generation of low latency infrastructures will be ubiquitous adoption of industry standards such as FIX Adapted for STreaming (FAST) and a focus on tuning the whole organism that supports low latency. This will include taking advantage of technological advances such as multi-core processors and acceleration tools that can be embedded in the silicon, such as Intel's I/O acceleration technology and its forthcoming XML and C++ library accelerators and specialized desktops.

Another important concept that will continue to make a significant contribution to the industry's low latency efforts is proximity hosting – positioning trading systems as close as possible to liquidity centres to shave off valuable milliseconds from execution times. Led by providers such as BT Radianz, which provides its Proximity Solution in the US, Europe and Asia, this practice, which has already been taken up by many banks, will grow, and firms will certainly continue to look at where they carry out functions today, and ask whether they can reap latency gains by moving them elsewhere. The entry of major operations into what is effectively new market practice represents an important hedge to the increasing potential operational and systemic risk of rapidly growing, highly tuned operations on both the sell and buy sides.

BT Radianz is among the several companies with which Intel is working in its Low Latency Lab facility – an environment offering banks and providers of technologies which can support low latency the opportunity to optimise solutions on state of the art hardware and infrastructure software, powered by Intel's latest Core2DUO chips, prior to moving them into production. As this activity grows, Intel can further emulate the real world by virtualising this capability between London, New York and the West Coast of the US and adding contemporary infrastructure software to work in various combinations in search of "optimum" - overall, providing a significant resource to support the industry's low latency optimisation efforts.

II. Introduction from A-Team Group

This research is the latest from A-Team Group (www.a-teamgroup.com) and focuses on the (mostly software) building blocks that combine to create low latency architectures.

A-Team Group continues to track the low latency technology space, and provide our own interpretation of events, through our *Market Data Insight* and *Transaction Networks & Technologies* publications.

We also recently launched an online information and news resource focused on this subject, at www.low-latency.com. We urge everyone to visit and make use of it, and to provide your own contributions and feedback.

We thank Intel Corporation for its support of this report and of our online presence, and for its focus on what we believe to be a fundamental enabler of the financial marketplaces 2.0 era. In the foreword to the report, you can read Intel's own thoughts on the subject.

As we explore in this report, latency is a given in all trading systems. It has always existed, and under the current known laws of physics, always will. But in recent years, several different technology approaches have been developed to minimize latency to the point where milliseconds are significant.

That milliseconds are significant is important, because milliseconds matter to the trading applications and market structures in place today. In the Market Landscape and Technology Response section, we examine in more detail the business drivers that have led to the imperative to implement low latency architectures.

This report focuses mainly on technologies that can be deployed within end user firms (and also marketplaces) to reduce latency for information delivery and transactions. We have sought to categorise these technologies into a number of functional silos, and in the Low Latency Architecture Technologies section there is an explanation of these silos, and the functional characteristics of products that fit into them.

But, as a preview, the silos that this report is focused on comprise:

- Datafeed Handlers.
- Messaging Middleware.
- Data Fabrics (or Data Grids).
- Event Stream Processing Engines.

We have profiled a number of technology vendors offering products in the above categories, and sought to document pertinent information from them. We have preceded all the above with our own suggestions for questions that might be asked to determine the best technology approach and vendors to solve your own business problems.

What is clear, however, is that the low latency technology world is evolving, and rapidly so in some areas. To this end, A-Team offers a few of our own predictions and thoughts on what is ahead for this segment.

1. With algorithmic trading spreading geographically and into new asset classes, increasing market data volumes and fragmented liquidity as drivers, investment in low latency architecture will increase substantially during 2007 to 2009.
2. There will likely be partnerships and some consolidation driven by a wish to provide more complete functionality and increased geographical presence. Only in extreme, specific cases are we likely to see consolidation through business failure.
3. Basic functions such as datafeed handlers, in-memory databases, messaging middleware and data fabrics will become increasingly commoditised. Vendors will enrich products with higher food chain functionality to sustain and grow margins.
4. Vendors of essentially toolkit products will provide increasing levels of professional services, perhaps via partnerships.
5. We should expect further acquisitions of specialised low latency players by large enterprise software players, e.g. Oracle/TimesTen and Progress/Apama, and maybe by larger financial applications vendors.
6. Benchmarking initiatives will likely be confined to individual vendor-led trials or small closed user groups. Some best practices may result. But no commonly accepted industry-accepted benchmark set will evolve.
7. Custom hardware approaches, e.g. FPGA and co-processors, will remain niche solutions. Multi-core mainstream CPU approaches, and higher speed networking, will be the main drivers for increased hardware performance.
8. Software products will increasingly make use of n-core chipsets and be grid-enabled.
9. Open source software projects will address at least some of the functionality that makes up low latency architectures and will be adopted.
10. Proximity Services will begin to offer higher value services, beyond co-location.

We look forward to further developments in the low latency space, and to seeing whether our predictions pan out.

A-Team Group
March 2007

SECTION 3: Vendor Profiles

IV. Event Stream Processing Engines

Event Stream Processing systems process incoming data messages, providing some value-added functionality, such as performing calculations.

Things to consider:

- Ease of use and functionality of APIs.
- Whether some datafeeds are bundled (or available) as standard.
- What level of data caching and replay is supported.
- What calculations are supported as standard.
- Ease of integration of custom calculations.
- How data is persisted to disk, requirement for third-party database, and performance implications.
- Availability of benchmarks.

v. Kx Systems

Headquarters location	Palo Alto, CA
Date founded	1993
Product	Kdb+
Number of clients	Kx has thousands of user licences in production. The approximate geographic breakdown is 60 per cent US and 40 per cent Europe.
Reference clients	Fidelity, UniCredit HypoVereinsbank (HVB), Lehman Brothers, Deutsche Bank and JPMorgan.
Pricing model	Kx solutions are priced on a per-CPU basis.
Example pricing	Depends on customer requirements and internal infrastructure. Pricing is available upon request.
Contact details:	Tel. +1 650 798 5155 Web. www.kx.com

Claimed Key Differentiators

Kx offers an integrated architecture that enables all components of the kdb+ platform to work together with optimal efficiency. The integrated kdb+ platform offers:

A relational database management system (kdb+).

A high-level language, "q" used for querying, programming, analytics and event management.

A built-in web server, file operations, communications and administrative tools.

Kx provides a unified solution for processing real-time, streaming and historical data within the same architecture. The data in the streaming, real-time and historical databases is saved in identical formats. This means a developer can write a query once and it will be able to run unchanged against streaming, real-time and historical data.

Datafeed Support

The kdb+tick add-on product is bundled with handlers for Reuters, Tibco and Bloomberg feeds; other data feed handlers are available on request.

Kx provides interfaces to common data feeds and it makes the source code for the feed handlers

available for clients who want to create their own.

First Derivatives, a Kx partner, can also implement custom feed handlers.

Middleware Support

Wombat and other middleware packages work with kdb+, but Kx says most of its customers use their Kx product to handle complex event data directly, to reduce latency.

APIs/Data Access

APIs supported for receiving and publishing events are C++, Java and .NET. Most Kx customers performing complex event processing work entirely within the kdb+ environment to reduce latency.

APIs are the same regardless of the data being handled (streaming, in-memory, derived, historical/time-series), because of Kx's integrated architecture and the fact that all data is handled using the same built-in array language.

Kx claims there is no need to marshal or unmarshal the data to accommodate different data representations. Kx supports time-series replay via kdb+tow.

Calculations Provided as Standard

Common calculations provided include VWAP, high/low open close (over any time bucket), and any kind of aggregation (including VWAPs) by user-defined time buckets.

Client-Specific Calculations

Kdb+ is customizable using its own language or by embedding existing C routines.

Clients have the choice of doing their own in-house development using Kx's integrated language, or they can hire outside kdb+ consultants if they prefer.

Studio for kdb+ is available as freeware and provides an IDE for customers preferring to work in a visual programming environment.

In addition, customers use Kx's C++, Java and .NET interfaces to write their own front ends.

Kdb+ and its layered applications also support ODBC and JDBC packages for access to the kdb+ database from other databases. For user-friendly output, most customers use Microsoft Excel, Matlab and other industry standard display packages.

Persistence

Persistence is handled entirely within the kdb+ product, which includes a historical database. Kdb+ also supports any third-party database using ODBC or JDBC, as needed by the customer.

Scalability

Nearly all Kx customers performing complex event processing see linear scalability as CPUs and

storage are added. Multi-threading and multi-core hardware and co-processors are fully supported and the product is being deployed in grids (in academic research settings). However, Kx claims the high speed and very low latency of kdb+ make grids unnecessary in customer deployments.

Resilience

Resilience and synchronization across multiple installations are achieved via replication of multiple servers. There is no limit to the number of servers supported.

Recommended Hardware/OS Configuration

Hardware: Intel, AMD and Sparc servers.

Operating systems: Windows, Linux and Solaris. For complex event processing, Kx strongly recommends 64-bit versions of the O/S being used.

Performance

Complex event performance: over a million messages captured per second on any hardware or operating system.

Embedding Within Other Applications

Kdb+ is running embedded in numerous custom applications. Embedding is accomplished via open, simple APIs.

Enhancements Planned

Kx releases new versions of its core kdb+ database approximately every six months.

Additional Information

Kdb+tick is an application that layers on top of the kdb+ database to capture, manage and analyze streaming ticks as well as terabytes of historical tick data. Kdb+tick can handle a million updates/second.

Kdb+taq is an application that layers on top of the kdb+ database to load NYSE TAQ data into kdb+ for relational analysis and real-time results on time-series analysis. It supports all NYSE TAQ formats - FTP/CD/DVD - and can load one day's worth of data in one minute.

Kdb+tow is an application that layers on top of the kdb+ database enabling traders to execute real-time analysis while replaying historical data to test algorithms for their portfolio models. It can replay a day's worth of market data in less than one minute.

Kdb+x Exchange Data Loaders are a family of stock exchange data loaders built for specific exchange products worldwide.