

Low Latency 5 Nine's Trading: Sun's Value Proposition for Market Data



The financial markets in 2007 are facing an unprecedented number of market wrenching challenges. Two of the key challenges are the continued growth of algorithmic trading and the impending enactment of trading regulations particularly in the U.S. and EU.

Algorithmic trading is driving firms to re-architect their trading systems for low latency and performance. Firms are automating trade processes and adding throughput and processing power to their market data and trading systems.

In 2007, Regulation NMS (Reg NMS) in the U.S., and the Markets in Financial Instruments Directive (MiFID) in the EU, will require most broker-dealers to seek best execution for their clients. In response, firms are adding direct market access (DMA) to exchanges, optimizing their order routing and execution, putting in place high availability systems, and doubling their data storage. Firms that fail compliance face loss of order flow and/or regulatory fines.

Why Market Data

Sun is closely working with leading application partners to address the latency and performance issues across the end-to-end trade process. A key focus is the market data space. The annual doubling in market data volume over past years will likely be eclipsed in 2007 due to Reg NMS and MiFID. Older market data systems are unlikely to keep up (as evidenced by a publicized hiccup at a major exchange), and trading firms are moving quickly to upgrade these systems. Sun and partners are focusing on how financial institutions can:

- Scale and process messages at a rate that may surpass 750,000 messages per second in 2007

- Find and reduce latency within and between market data and trading systems with Solaris™ 10 features like DTrace for observability and Containers for application consolidation
- Minimize operating risk including the risk of catastrophic failure and security breaches
- Implement these changes cost-effectively and with minimal interruptions

And unlike other mission critical applications such as CRM, as market data volumes continue to expand due to the introduction of new instruments and to business expansion, firms will likely refresh and expand their market data and associated systems by increasing processing and capacity capabilities on an ongoing basis.

Why Sun for Low Latency, 5 Nine's Trading

Sun's value proposition is unique. Through the combination of partnerships with the world's leading application vendor solutions and cutting-edge technologies, Sun enables firms to achieve low latency, high availability trading.

Partnerships with Leaders in Market Data and Trading Technologies

In the case of market data, Sun's partnerships with organisations like Reuters and Wombat have demonstrated performances of millions of updates per second at low milliseconds latency on Solaris 10. For the latest benchmark figures please visit sun.com/third-party/global.

Sun has recently announced support of Intel processors and, as a result, will soon provide SPARC®, AMD, and Intel-based systems. Sun and partners are focusing on providing versions of market data and trading technologies for the Solaris/Intel platform. For availability information, please contact your local Sun representative or visit sun.com

Cutting Edge Technologies

Sun technologies, specifically Solaris 10, provide immediate and visible benefits to customers re-architecting their systems for low latency 99.999% uptime trading. While this paper discusses the advantages of Solaris 10 (versus Linux) in the context of creating an optimized market data system, the same value propositions can be applied in other areas of the trade process to create a low latency, high-availability trading environment (shown in Figure 1).

There are several key factors that differentiate Sun in the Market Data environment. The following section focuses on the advantages offered by Solaris.

1. Installed Base

Over 60% of the global RMDS installations are on SPARC on Solaris 8, 9 and 10. The SPARC architecture has been in use for trading systems for over a decade. It is well understood, its behaviors are predictable, its usage characteristics are well defined, and its uptime statistics are unmatched. For trading environments not concerned with every last millisecond of latency, SPARC is the safe choice. Solaris 10 SPARC provides enough RMDS throughput that is sufficient for a large set of customers, hence we see many of

Figure 1

Decision Analysis	Pre-Trade Compliance	Trade Order and Execution
Risk and P/L Analyses	Investment Policy Covenants	Order Routing
Trade and Scenario Analyses	Regulatory Rules	Trade Capture, Mechanisms, and Fill
Allocation Analysis (Sector, country, etc.)	Internal Rules and Standards	Position Calculation/Reporting
FIX communication between applications		
Direct Market Access (DMA) to CNs/ATs/Exchanges		

them staying on SPARC for RMDS 6.0. Solaris 10 makes commodity AMD and Intel based hardware available for Market Data applications, offering the best latency and throughput.

2. Containers

The Containers technology is unique to Solaris 10. It is available on all hardware platforms that Solaris 10 is supported on, which include a large number of SPARC, Intel and AMD systems. The use of Containers enables the consolidation of in-house trading applications with associated market data components, thus reducing latency by eliminating data “hops” between servers. As an example, Reuters RMDS P2PS and a customers’ trading application could be deployed on two Containers on a single system, with guaranteed

physical resources for both applications. Another example is using Containers to consolidate multiple feeds using Wombat. Sun has tested consolidating a large number of OPRA feeds on a Sun Fire™ 4600 box using Containers, with near linear scale. With the OPRA traffic predicted to reach 750,000 messages per second by the end of 2007, this solution significantly reduces the hardware footprint required for OPRA. The results of this scaling exercise are available.

Hence Containers can be viewed as a way of consolidating applications to reduce latency between them, and also reduce the hardware footprint, thus lowering data center costs.

3. Observability using DTrace

DTrace allows unprecedented visibility into the market data applications and the Solaris kernel, clearly detailing the interaction taking place between the two. We have created custom DTrace scripts for RMDS and Wombat which will allow our customers to identify latency bottlenecks in their Reuters or Wombat applications while running in production. They can also leverage DTrace scripts to find bottlenecks in their in-house applications, provided the applications run on Solaris 10.

DTrace has been instrumental in identifying bottlenecks not only in RMDS and Wombat on Solaris, but also in complete trading environments.

The combination of DTrace and Containers is very powerful, allowing one to follow the lifecycle of a trade as it enters the front end trading systems, through the messaging bus, to the order routing systems, to the exchanges and all the way back to the front end systems. Running the various components of the trading environment in Containers allows one to use DTrace and other Solaris debugging tools to follow the flow of the transaction.

In today’s environment where latencies are measured in microseconds, DTrace enables firms to measurably shorten the trade lifecycle by quickly identifying bottlenecks.

4. Performance

- Networking Enhancements (FireEngine and Yosemite) – Solaris 10 introduces FireEngine, a completely redone TCP/IP stack, entirely eliminating the SVR4 STREAMS based

architecture. TCP and IP have been merged, connections are now processed on the same CPU for their lifetime—resulting in better cache locality, and the entire stack has been architected with a lockless design. This has resulted in tremendously improved TCP performance on Solaris 10. Update 2 of Solaris 10 introduces similar work for UDP/IP—Yosemite, again resulting in dramatically improved UDP performance for Solaris 10. Relevant benchmarks for RMDS, Wombat and TIBCO highlight these improvements.

- Multiple Page Size Support (MPSS)– Solaris 10 includes MPSS, which increases the application efficiency for accessing large virtual memory segments. MPSS allows Solaris to dynamically switch the default page size used for the program heap, stack and anonymous pages. This reduces the pressure on fixed resources such as the hardware TLB in the CPUs, significantly reducing the TLB misses. We have observed a 5% performance increase for RMDS using large pages.
- Scheduling Classes – Solaris has multiple Scheduling Classes available for applications to run in. The Solaris kernel is fully pre-emptable, which is unique for a general purpose OS. Time Sharing (TS) is the default class. The other classes of interest in the market data space are Real Time (RT) and FiXed priority (FX). Both RT and FX allow processes to be fixed at a certain priority level, not penalizing them for utilizing their CPU allocation quanta. We have observed performance improvements in RMDS and Wombat using the FX scheduling class. In the case of RMDS we have seen a

7% performance gain on Solaris 10 Update 2 by assigning its processes to the FX class. Wombat shows similar results. Perhaps more significantly, using the FX or RT class eliminates jitter in the response time of the application.

Using RT shows similar behavior, however rogue processes running in the RT class have the ability to make a system completely unresponsive, because the RT priority range maps higher than the System class (which is where most OS related processes run). Solaris is unique in offering the FX class which obviates the need for running a RT process in production

- Memory Placement Optimization – MPO was introduced in Solaris 8 to mitigate the latency effects of NUMA systems. Most systems being delivered today—SPARC and AMD—exhibit NUMA behavior. MPO introduces the concept of Latency Groups, which are pools of resources “physically close” to each other. An example is memory DIMMs and CPUs on a single system board. Another example is the CPU and its associated memory DIMMs connected via the on-chip memory controller on an AMD CPU. MPO is on by default, and both RMDS and Wombat benefit from it.

Whenever possible, Solaris schedules threads on their “home” latency group which allows threads to access warm caches and severely reduces the paging overhead. This allows better scaling of applications on systems with more than 1 socket.

5. Reliability

Solaris is well known in the industry for its uptime characteristics. Solaris 10 has added capabilities to further enhance this capability. The Fault Management Architecture allows Solaris to pro-actively diagnose and root-cause faulty components, allowing the system to stay up even in case of catastrophic failures. This capability exists on SPARC and AMD based systems, and is being worked on for Intel based systems. The Service Management Facility simplifies the system administration model, significantly reducing human error. Sun is working with Reuters to convert RMDS into a Solaris Service, which will allow management of RMDS through standard Solaris SMF commands. Market Data systems number in the 100’s at typical large sites. Sun Connection enables provisioning and change management of all these servers simultaneously, easing the burden of patch and configuration management.

6. Security

Solaris is one of the most secure commercial operating systems on the planet. Market Data systems typically sit at the edge of a corporate network. Features such as Containers (which allow market data applications to run with reduced privileges), Secure Execution (which ensures that Solaris executes only digitally signed binaries), File Integrity checking (to ensure that all relevant files in Solaris have not been spoofed), and Secure By Default (which allows you to harden a system with a single command) further secure market data systems.

Summary

To achieve the low latency and optimized trading efficiencies required in the algorithm-driven, post-Reg NMS/MiFID trading environment, firms need the unbeatable combination of the newest trading and market data systems running on Sun’s SPARC and/or x86 platforms. With Solaris 10, partners like Wombat and Reuters are demonstrating unprecedented scalability, reduced latencies, and greater ease of support with their market data products. These same technologies can be applied in trading environments cost-effectively, reliably, and securely in one or multiple locations worldwide to create a truly low latency, 5 Nine’s trading ecosystem.