

RTI Data Distribution Service

Ultra Low Latency Messaging for Financial Services

Free Trial Download
at www.rti.com

“ATD conducted extensive in-house testing of the leading messaging products targeted at market data distribution. We found that RTI provided the most consistent performance, with no latency spikes even under times of peak volume. RTI also had the most full-featured product, with more comprehensive control over performance tuning.”

RTI Data Distribution Service provides a peer-to-peer, completely decentralized messaging infrastructure for distributed real-time applications. It is the industry’s highest performance, most reliable, flexible and intelligent messaging middleware.

RTI Data Distribution Service is battle-tested technology proven in hundreds of time- and mission-critical applications, including:

- Financial market data distribution, transaction processing and compliance
- Combat systems
- Air traffic control
- Medical equipment
- Industrial automation

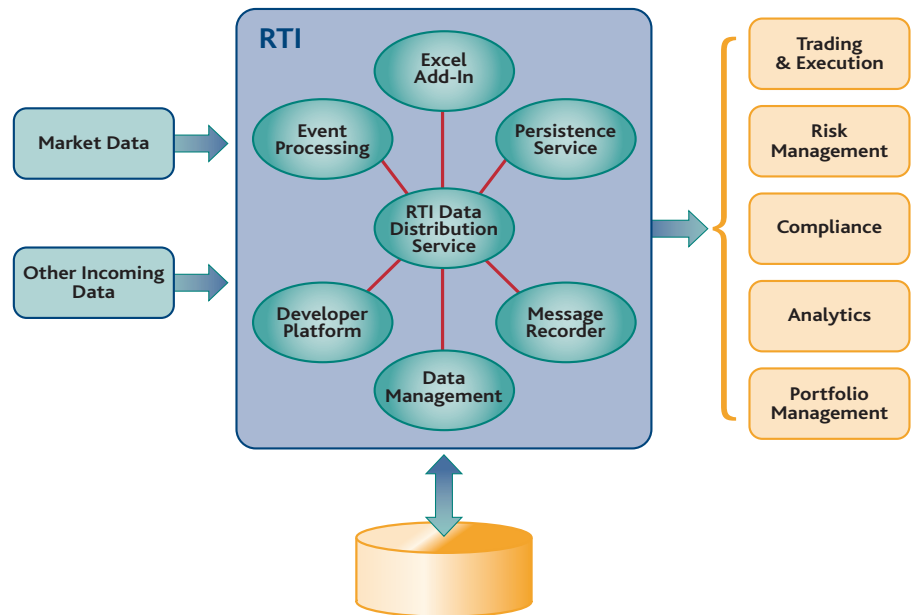
RTI Data Distribution Service provides seamless interoperability across a broad range of operating systems and programming languages, allowing deployment in highly heterogeneous systems.

High-Performance Messaging

The core of RTI Data Distribution Service is a high performance messaging layer combined with a distributed in-memory data caching function.

Leveraging Gigabit Ethernet and standard hardware, RTI can deliver millions of messages per second and features median latency as low as 30 microseconds. Aggregate network traffic is bounded only by the network switching fabric.

RTI Data Distribution Service features peer-to-peer communications, reliable multicast transmission and publish-subscribe semantics. Combined, these features provide for unmatched levels of determinism and scalability in message communication.



RTI Data Distribution Service is the core component of RTI's comprehensive low-latency infrastructure.



Highlights of RTI Data Distribution Service include:

Peer-to-Peer Communications—ideally suited to the transport of financial market data and transactions

- Messages are transmitted directly from sender to receiver for minimum latency, without routing via intermediate brokers, servers or “daemon” processes.
- Throughput is maximized, allowing full leverage of network switching infrastructure with no bottlenecks caused by server overhead.
- There are no single points of failure for high resilience.

Reliable Multicast—the most efficient method of delivering high-volume market data to trading applications

- One-to-many simultaneous message transmission—just one message is sent regardless of the number of recipients, delivering massive scalability.
- Message filtering is performed at the network switch level, removing processing overhead from servers.

- Messages can be balanced across multiple multicast groups—each node only receives messages sent to multicast groups corresponding to subscribed message subjects.
- RTI’s highly-tunable reliability protocol safeguards against dropped packets, network congestion and NACK storms.

Tunable Message Characteristics—balance latency, throughput and resource consumption to maximize performance by delivering data only where needed, not over the entire network.

- Quality of Service (QoS) properties (such as reliability, durability, delivery order, use of multicast) are independently configurable for each message subject, producer and consumer.
- Configurable batching of messages within network-level packets (including time and byte limits) minimizes processor overhead and maximizes throughput.
- Tunable traffic shaping sets limits on how much bandwidth a sender can consume, and how often messages may be sent.

This allows very large messages to be sent without impacting other senders’ latency and throughput.

High Determinism—consistent and reliable performance as market rates fluctuate

- Latency is consistent as message rates increase: variation from the linear scale is minimal.
- With a single producing thread, median latency for a 200-byte message is 79 microseconds at 100,000 messages per second. Latency grows to only 129 microseconds, even near network saturation at 525,000 messages per second.

Publish-Subscribe Messaging—efficient one-to-many asynchronous communications

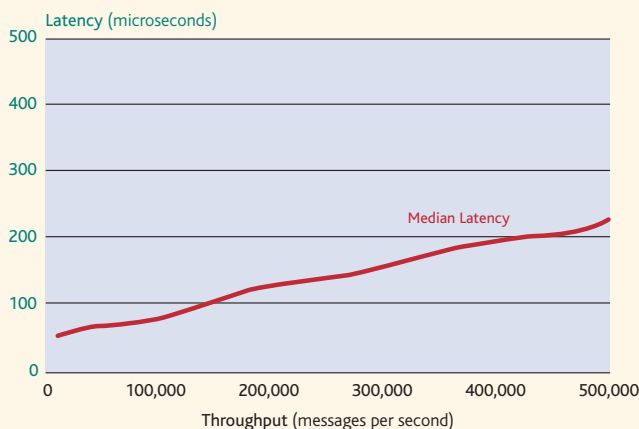
- Message definitions, such as field names and types, are exchanged once, at application start-up time. Only data is sent with each message. This increases performance and reduces bandwidth consumption compared to the conventional approach, in which each message is self-describing and thus includes a substantial amount of metadata in addition to actual data.

Profiting From Determinism

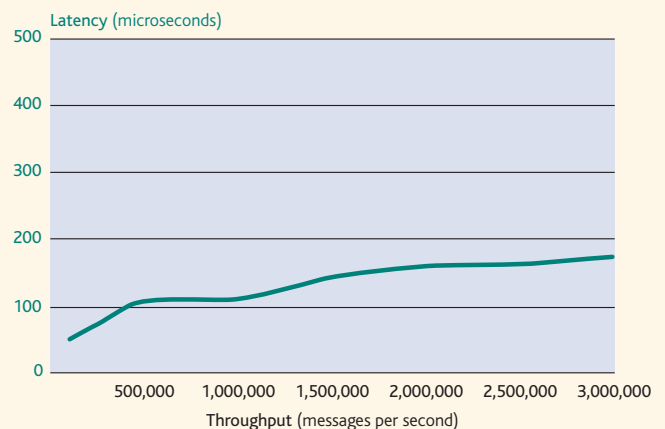
Opportunities to trade profitably are maximized as market volatility increases. But when volatility and message rates peak, legacy systems melt down in terms of consistency of latency, introducing ‘jitter’ which renders them useless for supporting many financial markets applications.

Algorithms and smart order routing systems that exploit volatile markets through statistical arbitrage and other models rely on a consistent latency that can be built into the algorithm’s logic. Thus, determinism of latency is essential for high-frequency trading applications.

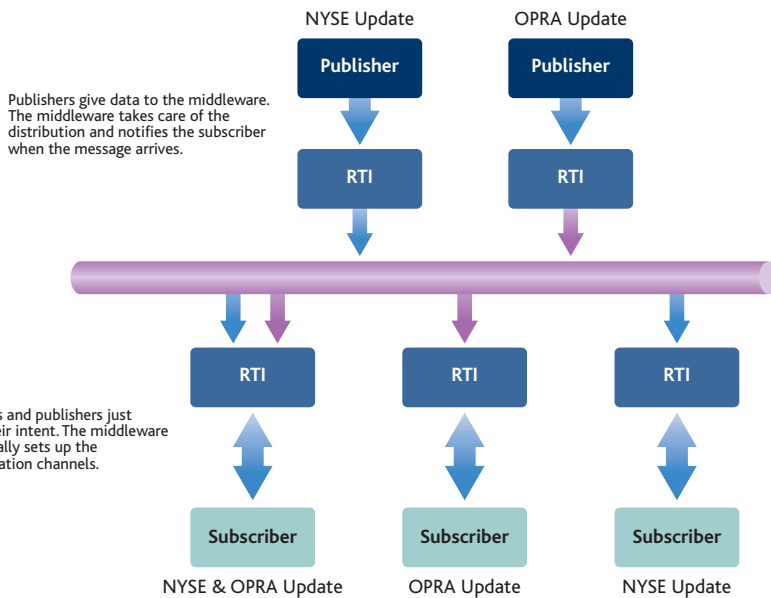
Impact of Throughput on Latency for 200-Byte Messages
Single Publishing Thread and NIC



Impact of Throughput on Latency for 200-Byte Messages
Eight Publishing Threads and NICs



These benchmarks were conducted with reliable multicast over Gigabit Ethernet between machines with 2.4 GHz Intel Core 2 Duo E6600 processors.



RTI's model takes care of channel configuration and data distribution for the applications.

- Recipients subscribe to subjects such as messages from a particular market, particular symbol or a combination, such as “Trade prices for IBM common stock listed on the NYSE.”
- Subscriptions may include wildcards representing a set of subjects, and can subscribe to hundreds of thousands of subjects (e.g. all stocks listed on a market) without a significant performance impact.
- Messages can be filtered and subscribed based on content, not just subject.
- Slow consumers—recipients unable to deal with the flow of messages to them—can be easily handled in a number of ways. These include message conflation, alerts when buffer watermarks are reached, message loss notification, and slow subscriber bypass and opt-out.
- Publishers are notified of new subscribers and may reject those not authorized to receive certain messages (e.g. from a particular market).

Persistence—RTI Persistence Service provides message replay and high availability even if the original message producer is no longer running.

- The RTI Persistence Service can receive messages in parallel with other consumers so as to not impact latency.

- Durable subscribers that go offline receive missed messages when they rejoin the system. New subscribers may also request a replay of prior messages.
- Messages may be stored in an in-memory and/or on-disk persistence store for the optimal balance of latency, data durability and history.

- The persistence service can be used as a broker for subscribers that do not require minimum latency. This allows publisher resources to be devoted to those subscribers that require the lowest latency.
- Multiple persistence stores can run concurrently for redundancy and load balancing.

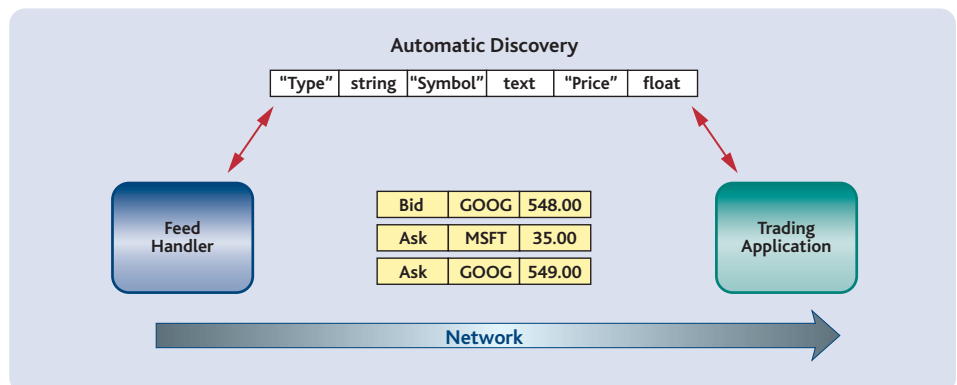
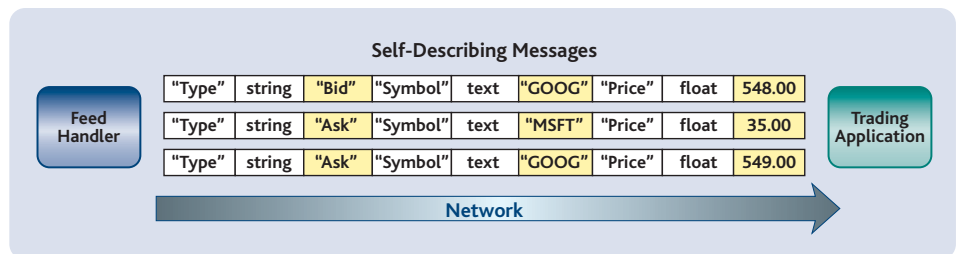
Fault Tolerance and Redundancy—architected for high availability

- Messages may be simultaneously transmitted via multiple network interfaces.
- Redundant publishers of the same messages can be configured, with automatic failover.
- All services such as recording, persistence and database integration are redundant, providing high availability.

Distributed Caching

Distributed in-memory data caching is built directly into RTI’s messaging layer allowing applications instantaneous access to data objects.

The cache can be used for real-time data, such as best bid and offered prices. In addition, with RTI Real-Time Connect, the cache



With RTI, applications can exchange message definitions at startup (discovery) time. This provides the benefits of structured and typed messages without the network overhead incurred by traditional self-describing messages.

RTI Data Distribution Service product information

can be synchronized with a database, e.g., a ticker plant. This provides SQL applications with access to real-time data and allows real-time applications to quickly access enterprise data, such as market pulse and time-and-sales, via the local cache.

Highlights of RTI's caching capabilities include:

Integrated with Messaging—eases application design and improves performance

- Applications that do not require a continuous stream of data can request the last value from the cache and receive it instantaneously.
- Publishers are shielded from details of whether messages are being consumed in real-time or as requests from the cache.
- Performance is increased versus separate messaging middleware and data caching products.

Configurable History—provides built-in time series support and enhances application performance

- Whether data objects are stored on a last value basis or as a time series is configurable by data object.
- The lifespan of data objects can be set so they are automatically deleted after a certain time, ensuring applications do not retrieve stale data.

Backend Database Integration—with RTI Real-Time Connect, real-time applications store and retrieve data from a database

- Caching layer provides instantaneous access to frequently retrieved objects stored in the backend database.
- Support is included for Oracle Database, Oracle TimesTen In-Memory Database and MySQL.

Representative Customers

RTI's financial service customers include:

Automated Trading Desk, LLC (ATD), a part of Citigroup, is using RTI's real-time messaging infrastructure as the foundation of its market-data distribution infrastructure. ATD is using RTI's middleware to distribute real-time data from direct-exchange and Electronic Communication Network (ECN) feeds to its price-prediction engines, automated trading applications and order management system (OMS).

PIMCO, one of the world's leading fixed-income managers, has selected RTI's high-performance messaging middleware to expand the functionality and enhance the scalability of PIMCO's compliance-checking software.

Zivlyn Systems, a leading provider of next generation distribution systems for financial institutions, selected RTI for its next-generation trading system and order book system.

Specifications

Platforms

- Linux 2.4 and later
- Solaris 2.8 and later (SPARC and 2.10 (x86))
- Windows 2000 and later
- Windows CE/Mobile

Processor Families

- AMD 64/EM64T
- SPARC (32 & 64 BIT)
- X86

Transports

- IPv4
- IPv6
- Shared Memory
- Pluggable Interface
- InfiniBand

Language Support

- ANSI C (C89) and ISO C (C99) with GNU extensions
- ANSI C++
- Java 1.4 or later
- Real-Time Java (RTSJ)
- .NET

Standards Compliance

- Object Management Group (OMG) Data Distribution Service for Real-Time Systems (DDS)
- SQL

Note: RTI continually adds support for new platforms. Please contact RTI for additional availability.

About RTI

Real-Time Innovations (RTI) is a leading provider of extreme performance messaging, real-time data management and complex event processing software.

RTI's technology has been proven in over 400 mission-critical applications in financial services, medical, transportation, communications, industrial control, aerospace and defense. Customers include PIMCO, Automated Trading Desk, Varian, Nikon, Schneider, Lockheed Martin and Raytheon.

Founded by Stanford University scientists in 1991, RTI is privately held and headquartered in Silicon Valley.

US HEADQUARTERS

Real-Time Innovations, Inc.
385 Moffett Park Drive
Sunnyvale, CA 94089
Tel: (408) 990-7400
info@rti.com

OFFICES

New York
1 Liberty Plaza, 23rd Floor
New York, NY 10006 USA
Tel: +1 (212) 201-2910

United Kingdom

Gloucester House, 399 Silbury Blvd.
Central Milton Keynes, MK9 2HL
Tel: +44 (0) 1908 303 685

www.rti.com

