

Accelerating and Simplifying Data Distribution for Exchanges

Exchanges and Alternative Trading Systems are inherently high-volume, require the lowest possible latency for competitive differentiation, and demand high availability due to the far-reaching ramifications of slowdowns or outages.

Most liquidity providers build their systems using a multi-tiered architecture that consists of order entry systems and FIX gateways to bring in and normalize orders, matching engines to create trades by matching bids and asks, and feed gateways to distribute results and updated pricing information to market participants. They use a similarly complex set of disparate messaging middleware products to route information within and between these disparate environments because of the need for many qualities of service such as low latency, persistent, IPC and WAN.

Solace provides a unique hardware-based solution that supports all of the qualities of service exchanges need with best-of-breed performance in each area. Solace message routers are easy to deploy, operate, upgrade and scale—all of which contributes to lower TCO than software-based solutions.


This paper describes the areas within exchange systems where Solace messaging products provide value, provides an overview of Solace's products and capabilities, and summarizes how some customers are using Solace today.



The Role of Middleware in Exchanges

The architecture of typical order matching systems is shown in the following figure.

- **Order Entry or FIX gateways** This layer faces sell-side clients and terminates the protocol from clients submitting orders to be executed, converts them to an internal format and coordinates execution with the appropriate matching engine for each security being traded.
- **Matching Engines.** This layer maintains the order book and handles all bid/ask orders to create trades. This layer is typically scaled by horizontally partitioning the books for various securities across instances of the match engines.
- **Feed Gateways.** This layer receives the aggregate output of the matching engines in the form of bids, asks, trades in internal format, creates external formats of these messages and then outputs this flow to the Street. This could be carved into different multicast groups, or fanned out over TCP. Also, the data could be filtered into different groupings to produce different feed products prior to distribution.



Connectivity between the various layers is typically established by messaging middleware. This middleware must be capable of high rate and low, consistent latency. It also must support both reliable and guaranteed messaging qualities of service, provide high availability and include a high degree of management visibility and control. These are precisely the characteristics that differentiate Solace messaging products from other messaging products.

The speed and predictability of information flow between order entry gateways and matching engines, as well as between matching engines and feed gateways, is critical to the success of any liquidity provider.

Messaging from Gateways to Matching Engines

From gateways to matching engines, the underlying messaging system must support high, volatile and unpredictable message rates. It also needs to selectively route order requests to the right matching engine for the security being traded, and route responses back to the originating gateway. To attract high-volume traders, exchanges must ensure consistently low latency from order entry to order acknowledgement back to the sell side at all times.

Messaging from Matching Engines to Feed Gateways

Messaging from matching engines to feed gateways also demands high rate with low latency. A guarantee of delivery is required too, though, because market participants must have an accurate and up-to-date view of all market activity to avoid basing decisions on an outdated understanding of the market.


Messaging to Downstream Systems

There are also less latency-sensitive applications that must monitor the order flow for market surveillance, risk management regulatory compliance, data warehousing, etc. It's critical that they receive all messages in sequence no matter what, and that they never impact the real-time communication between order entry gateways and matching engines.

Use Cases

Order Execution

By providing both reliable (in memory) and guaranteed (stored in secure storage) messaging in a single platform, Solace provides the ideal infrastructure for a system with such varying requirements. For a given message flow, some applications can have non-durable subscriptions so they can receive the messages with the lowest possible latency while other applications can use durable subscriptions so they are guaranteed to never lose a message even if they temporarily cannot keep up with the message flow. Solace supports both “store and forward” persistence (with sophisticated queue servicing capabilities) and “forward and store” persistence for ultra low latency. Class of persistence can be dictated by subscribers without any impact on the publishing application.




Global Sequencing

Some exchanges require that all subscribing applications (e.g. feed gateways) receive their messages in the exact same order, even in the presence of multiple uncoordinated publishers. Even with multicast messaging systems, solving this problem most efficiently requires a single sequencer to determine global ordering however in software-based systems, this sequencer quickly becomes both a performance bottleneck and an additional source of variable latency. A hardware-based message broker that naturally serializes parallel message streams at extremely high rates automatically ensures a common message order to all subscribers.


Real-time to Near-Real-time Buffer

All trading platforms need to feed the output of the real-time trading system to a variety of downstream applications such as trade data capture, settlements, market surveillance and more. Even in architectures where the three tiers described above are provided as a combined “black box” by a third-party vendor, with little opportunity to insert a different messaging system, trade data must be reliably captured and sent to the exchange’s many downstream custom applications without loss at a rate they can accept.




Solace message routers can meet that need by serving as a “high performance buffer” between the real-time systems and the near-real-time systems. Solace’s hardware essentially absorbs messages produced by the real-time systems, persists them without impacting the performance of the real-time systems, and queues them for subsequent fanout to near-real-time applications. This enables the efficient handling of slow consumers without impacting the performance of publishers.

Each message router can route over 450,000 guaranteed messages per second at very low, consistent latency even in the face of slow or offline consumers. This chart shows how egress bandwidth increases when messages of varying sizes are delivered to as many as 50 clients each.



TCP Feed Distribution

In some cases, exchanges distribute market data to market data customers over TCP. Feed gateways typically generate the feed data (including any heartbeats) and send that feed to a bank of servers which provide the TCP fanout to clients. Typical issues with this approach are the relatively small number of fanout clients per server thereby creating datacenter sprawl, the uncontrolled “first-to-last” latency and lack of management visibility into the performance of each connection.



Solace message routers perform TCP fanout in hardware, offering very tight first-to-last latencies thanks to its ability to transmit messages back-to-back on 10Gbps.

In addition, since the datapath is implemented in hardware, Solace message routers provide a rich amount of management information on a per-connection basis at every layer of the protocol stack. This includes TCP round trip times, TCP retransmitted packet counts, per-connection message rates (1sec and 60sec) and per-connection queue depth monitoring (current and maximum). This information, which is not made available by operating systems, is critical to diagnosing the source of poor performance over WAN connections. All of this is available from Solace message routers without any performance impact due to the natural parallelization of hardware.

Solace's Proven Solution

The Solace message router enables organizations to increase the speed of their application infrastructure by 10 to 100 times, improve the manageability of their system and reduce costs. It supports reliable messaging and non-persistent JMS, as well as guaranteed message delivery and persistent JMS. All of its capabilities and services are accessible and manageable through Solace's unified API and administration framework. Solace's hardware-based middleware performs messaging with unprecedented performance, reliability and manageability.

Types of Messaging Supported

- **Reliable messaging:** Solace's reliable messaging solution delivers low latency even at very high volume.
- **Guaranteed messaging:** Solace uses an innovative, patented approach to persistent message queuing that enables guaranteed messaging an order of magnitude faster than software-based alternatives.
- **WAN replication/synchronization:** Solace enables high-performance WAN distribution by sending information only where it's needed, compressing messages, and fanning out messages at the edge.
- **Web messaging:** Solace enables the quick, cost-effective deployment of dynamically-updated user interfaces with a turnkey message router that streams real-time data over the internet.
- **IPC Messaging:** For the most latency-critical scenarios where co-locating applications on a single server is possible, Solace 's API supports IPC messaging in a shared memory environment with average latency of 400 nanoseconds.

High-level Benefits

The Solace message router offers the following advantages and benefits:

- **Speed:** Solace eliminates the drawbacks of software running on general purpose servers to enable the delivery of massive amounts of information at higher rates and with less latency and less latency variation than other solutions.
- **Simplicity:** By handling all of the major types of messaging in a single platform, Solace enables companies to reduce the complexity and footprint of their infrastructure and accelerate time-to-market for new applications.
- **Savings:** With their high capacity, support for virtualization and ability to support many different messaging types on one device, Solace message routers enable dramatic consolidation.
- **Stability:** Since Solace's solution is implemented in hardware, it maintains its performance characteristics even when processing millions of messages per second and in the face of misbehaving client applications.
- **Scalability:** Solace message routers are self-sufficient network devices that work together in concert, so Solace provides the linear scalability people are used to getting out of their IP routers and switches.

Solace's hardware-based middleware performs messaging with unprecedented performance, reliability and manageability.

By handling all kinds of messaging in a single platform, Solace lets companies reduce the architectural complexity and physical footprint of their IT infrastructure.

Public Production Deployments



London Stock Exchange Group selected Solace's technology to serve as the persistent messaging and buffering layer between their real-time trading system and the downstream systems used for surveillance, data warehousing, post trade processing and billing.



The TMX Group chose Solace's solution as the messaging backbone within their data dissemination plant, the core technology that performs the distribution of Toronto Stock Exchange and TSX Venture Exchange data feeds to market participants.



Korea's sole securities exchange operator has selected Solace messaging message routers as the core data delivery infrastructure for their new high-speed exchange platform. The new system, called EXTURE+, is the KRX's next-generation trading system that offers low latency trading across many asset classes including equities, bonds, and derivatives.



The Canadian Securities Exchange uses Solace to distribute high speed data feeds to clients. CSE selected Solace's solution based on its ability to improve the company's infrastructure without client impact or application changes, enabling them to capitalize on the benefits of the new system quickly and with minimal risk.



Liquidnet is the premier institutional investment community, bringing together the world's largest asset managers and public companies on a single network that directly connects traders, portfolio managers, analysts and corporate issuers. Liquidnet selected Solace message routers as a key component of their liquidity matching service.

Capabilities and Performance

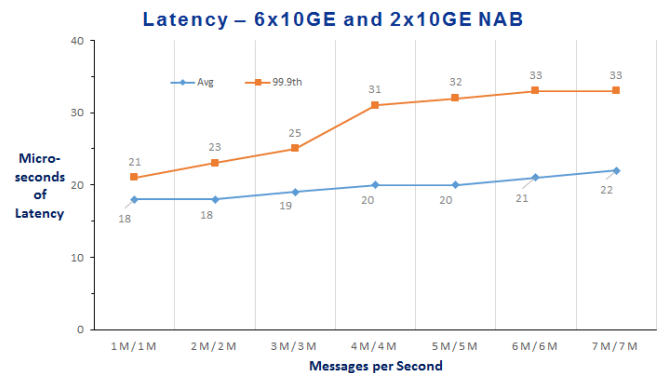
Reliable Messaging

Reliable messaging establishes a real-time flow of information between many applications and sources of information such as market data feed handlers, sensor readings or emergency alerts. Reliable messaging is appropriate for scenarios where achieving the lowest possible latency is more important than a guarantee that every single message be delivered every time, and for situations where the real-time flow of information is only relevant and wanted by subscribers who are currently connected.

By handling the routing of messages in purpose-built hardware, each Solace message router can deliver millions of messages per second with very low, predictable latency.

The performance results to the right reflect the delivery of 100 byte messages from publisher API to subscriber API. To achieve optimal latency, cut-through Ethernet switches were used with 10GigE links to stream packets and client kernel bypass was used to reduce latency on the client machine.

Solace's solution makes it easier to enable information flow across systems that consist of different types of devices connected via network links of varying throughput by buffering messages and playing them out at a rate each link and recipient can handle.




Guaranteed Messaging

There are many distributed application scenarios that can't tolerate any loss of messages—every message must be successfully delivered, in sequence, or reported as undelivered if a recipient becomes permanently unavailable. This guarantee can be handled by the sending application or messaging middleware—in either case a copy of each outbound message is stored until receipt is confirmed by all of the intended recipients. Handling the guarantee of message delivery in middleware greatly reduces application development time and complexity as well as CPU utilization, enabling greater performance and scalability.

With the ability to route 450,000 messages per second, and 1.6M messages per second in fanout scenarios, Solace message routers are 20-100 times faster than software-based alternatives. It also requires a much smaller footprint than software-based systems, reducing architectural complexity and cost of ownership.

Solace's solution is able to achieve this unprecedented performance by storing messages in on-board memory that survives power failures instead of persisting them to disk like conventional software-based solutions.



Web Messaging

The increasing demand for real-time data has companies seeking ways to stream information to users at their desks and on the go. Two developments are paving the way: Internet push/streaming technologies such as Comet and WebSocket, and web-centric application frameworks such as HTML5, JavaScript, Silverlight, .Net, Flash and mobile apps.


Solace's solution can extend messaging services for internet data streaming, cost-effectively providing ubiquitous reach to applications and users inside the organization and anywhere in the world via the Internet and mobile networks/devices. Solace makes messaging over the internet a secure extension of the internal messaging network rather than a custom integration/translation between two disparate technologies.



IPC Messaging

In the latency-sensitive scenarios such as high-frequency trading, companies are always looking for techniques that will squeeze latency out of their systems. One such technique is Inter-Process Communications (IPC), a way of sharing data between applications hosted on a single multi-core server.

By eliminating the latency of sending information across the network between computers, Solace supports IPC messaging with average latency under 400 nanoseconds. Trading applications inevitably have to pass information to applications running on other systems to complete and settle transactions. Since Solace's API and message routers support many kinds of messaging, applications can share information with both co-located applications and remote applications using other messaging types using the same API.



WAN Optimization

The amount of data traversing the globe via wide area networks (WANs) is surging, in many cases beyond the capacity of existing systems which causes slowdowns and unpredictability. Adding bandwidth is an effective (if expensive) solution, but only gets you so far before latency and application behavior create diminishing returns.

Solace's message routers can be used individually for fast fanout over the WAN, or as the foundation of a sophisticated multi-datacenter WAN optimization system. When data rates exceed capacity messages are buffered so the system can weather outages and traffic spikes without losing messages. Solace optimizes bandwidth by only sending messages across the WAN when an application on the other side is subscribed to it, by performing streaming compression on a per-client basis, and by offering per-client rate control.

Message Caching

It's important for some systems to cache data so clients can look up recent quote data to find the last value or investigate historical activity. Solace offers a fully integrated caching solution that can cache 500,000 messages per second, and perform 70,000 lookups per second. It can be run on a Solace device or standalone servers for load-balancing, redundancy and maximum scalability. Cache topics can be distributed so the solution can be easily scaled to meet increasing requirements for rate and topic space.

Solace's message routers can be used individually for fast fanout over the WAN, or as the foundation of a multi-datacenter WAN

Advantages of Solace's Unique Approach

Shareable Infrastructure

Solace supports the creation of "Message VPNs" that can be used to give many applications and departments their own fully compartmentalized virtual messaging environment within a single Solace message router.

When message VPNs are set up messages never cross between environments, and administrators can fully configure the behavior of, and resources available to each virtual messaging partition. This can reduce TCO and accelerate time to market by enabling many applications to tap into the power of a single easily-managed device. They can even use this functionality to run several discrete development and test environments in one box.

Messaging as a Service

Solace makes the routing and delivery of information available as a network service so developers can create applications without worrying about how they'll interact with users and other applications. They simply build the application to perform its intended function and tap into the Solace-based infrastructure for messaging functionality. This reduces time-to-market for new services and keeps application complexity to a minimum.

Unified API and Administration

All of Solace's features and capabilities are accessible through a single API that's available for all of the most common operating systems and programming languages. It's a robust, high performance, full-featured yet low footprint API very similar to the APIs for software-based middleware, so developers can leverage their existing skills. And since the one API is used for all kinds of messaging, companies can repurpose people, tools and templates across all of their messaging initiatives. Solace's administration framework makes it easy to monitor and manage a Solace-based system. It provides comprehensive per-client statistics without impacting performance at all, and can be integrated with existing monitoring and management tools.

Easy Flexibility and Scalability

Solace's solution provides flexibility and scalability on two levels. First, Solace message routers are chassis that can accommodate a number of PCIe cards (5 in a 3230 and 10 in a 3260) that perform functions such as protocol termination, message routing, guaranteed delivery, etc. This lets customers configure Solace message routers to meet their exact needs, and expand functionality over time within that same footprint. Second, Solace message routers propagate subscriptions and share information about network availability, so any and all of them can identify the shortest possible path for each message. This means the capacity and/or functionality of a system can instantly be increased by simply deploying another Solace device.

Client-specific Flow Control

Solace message routers supports per-client per-topic flow control for recipients who can't handle or don't need every update in a high-speed message flow, but do need to make sure they are always receiving the latest information. Solace addresses this requirement by consolidating updates to a topic on a client-specified interval and passing along only the

Solace makes it easy to give many applications their own secure messaging environments within a single message router.

Since a single API is used for all kinds of messaging, companies can repurpose people, tools and templates across all of their initiatives.

most recent update. Common use cases include streaming market data to human traders and rate limiting FX updates over a WAN.

Security

Since Solace's message router handles message distribution via discrete TCP connections, the platform can be made much more secure than multicast environments. Administrators and client applications alike must be authenticated using username/password in order to gain access to the message router. Solace supports an on-board authentication database, or integration with Radius and LDAP. Publisher and subscriber "Access Control Lists" determine what each application is allowed to send and receive, and IP layer ACLs restrict which IP address and subnet they can connect from.

Summary

Solace provides a unique hardware-based solution that supports all of the qualities of service exchanges need with best-of-breed performance in each area. The Solace message router is easy to deploy, operate, upgrade and scale—all of which contributes to lower TCO than software-based solutions.

To learn more visit
solace.com or call +1
613-271-1010.
