

The Advantages and Performance of Solace's Guaranteed Messaging Solution

There are many distributed application scenarios in which every message must be successfully delivered, in sequence, or reported as undeliverable. In fact, *most* interactions between computer systems require (or would prefer) a guarantee of delivery, hence the ubiquity of MQ and JMS (which offers persistence) in enterprise IT.

This guarantee can be handled by the sending application or messaging middleware, but in either case a copy of each message is stored until receipt is confirmed by intended recipients. Handling the guarantee of message delivery in middleware reduces application complexity as well as CPU utilization, enabling greater performance and scalability.

Solace's hardware-based guaranteed messaging solution is faster than software-based alternatives, and features less architectural complexity and lower cost of ownership.

This paper is intended to introduce IT professions to Solace's solution by describing how it works, providing detailed test results that document its performance, and explain how it differs from conventional solutions.



Introduction

Here are the primary advantages of Solace's solution:

- **Performance:** As documented later in this paper, each Solace message router can deliver over 450,000 messages per second, and fanout rates up to 1.6 million messages per second, all with low, consistent latency at a wide range of message rates and sizes.
- **Consolidation:** The limited throughput of software-based JMS and MQ message brokers has forced companies to deploy thousands of servers to achieve necessary capacity or performance. The combination of Solace's extremely high guaranteed messaging rates and its virtualization capabilities give each message router the ability to handle the workload of dozens of software brokers, enabling companies to significantly reduce the cost and complexity of their guaranteed messaging infrastructure while improving its robustness and scalability.
- **Fault Tolerance:** When a third of receivers were disconnected, there was no impact on senders, and mean latency for receivers that remained online was virtually unchanged. And in an aggressive "pull the plug" test, all messages were correctly routed and delivered in the proper order thanks to failover to the secondary device.

Scenarios that Require High-Performance Guaranteed Messaging

Solace's exceptional performance and resilience are especially beneficial in scenarios where high rate or fast, steady performance are a necessity. Here are scenarios in which high-speed, high-throughput guaranteed messaging can help companies achieve performance goals and dramatically reduce the TCO of their infrastructure through consolidation.

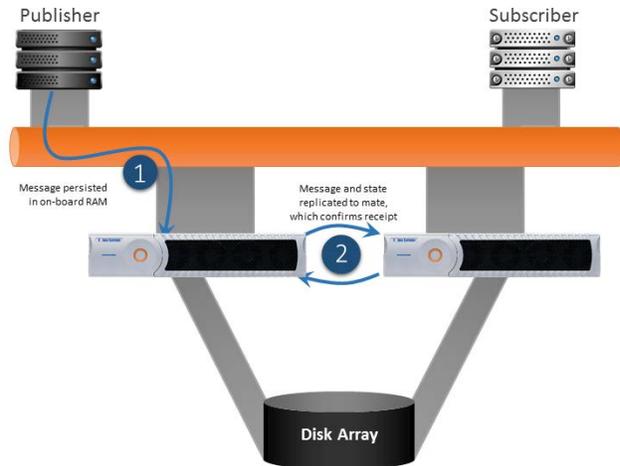
- **Trading Platforms** – Financial firms engaged in electronic trading need to ensure that orders sent to liquidity pools are never lost, and that responses to orders (acknowledgements, fills and partial fills) are never lost. This traffic is bursty and must be routed very quickly to ensure accurate fills and avoid slippage.
- **Execution Platforms** – Exchanges and liquidity providers need to move messages between gateways, matching engines, smart order routers and to execute trades, and all of that messaging must be guaranteed to satisfy clients and regulatory requirements. To learn about our solutions for exchanges, visit <http://solacesystems.com/exchanges>
- **Middle- and Back-Office Integration** – Within sell-side investment banks, trade events sent to and between back- and middle-office applications represent a handoff of responsibility from an intensely real-time system (the trading platform) to critical post trade applications (such as OATS reporting, trade data capture, settlements, intra-day risk) whose performance and capacity are much lower than that of the trading platform. Many customers need to transfer as many as a billion trade events per trading day into the middle office with full persistence, without ever losing a message, while fanning each event out 5 or more times and without ever slowing down the trading applications. Such use cases demand guaranteed messaging with exceptional capacity and consistent performance.
- **Real-Time Inter-Application Communications** – Many messages sent between enterprise applications represent a handoff of responsibility, so it's critical that messages are always delivered, and delivered in the same order they were sent. Examples include the execution of customer-facing business processes, delivery of eCommerce/Web portals, synchronization of transactional databases, and control systems in manufacturing and railways. In such situations the non-delivery or out-of-sequence delivery of messages can lead to bad decisions, dissatisfied customers and dangerous situations. In many cases companies have been forced to sacrifice a desire for real-time information exchange by living with the slow speed/throughput of software-based guaranteed messaging, and in others they've sacrificed failsafe delivery guarantee and accepted the risk of turning off failsafe mode and message monitoring because that was the only way to achieve latency or throughput goals.

Architecture and Operation

Solace uses a patented method to store messages in on-board memory instead of persisting them to disk. Messages are persisted to disk if on-board memory fills up, in which case messages are sent to disk in batches. The message router's high throughput capacity means slow consumers can reconnect at any time and "catch up" without impacting other subscribers.

Redundant Architecture

Solace message routers are deployed in an active/standby model to provide high availability. Messages are stored in high-speed on-board RAM on both message routers which are connected by two fiber optic links. In the case of a single message router failure, the back-up device takes over activity and allows applications to continue operation within five to ten seconds. If both message routers fail, there is an auxiliary power supply built-in that supplies enough power to flush the messages and their delivery state to an onboard flash disk so no data is ever lost. This data is read back into the message router hardware to resume service when the message router is powered up again.

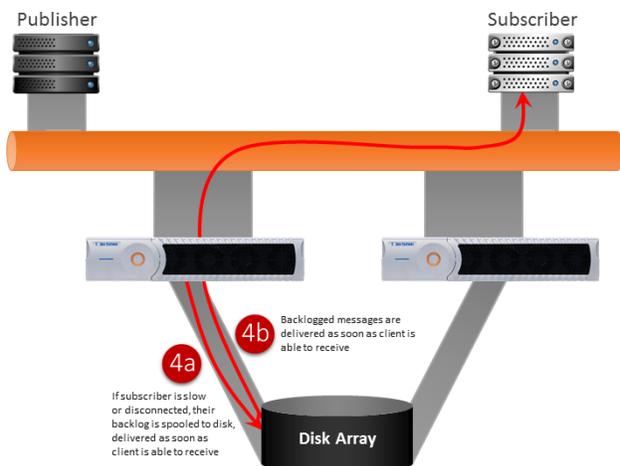
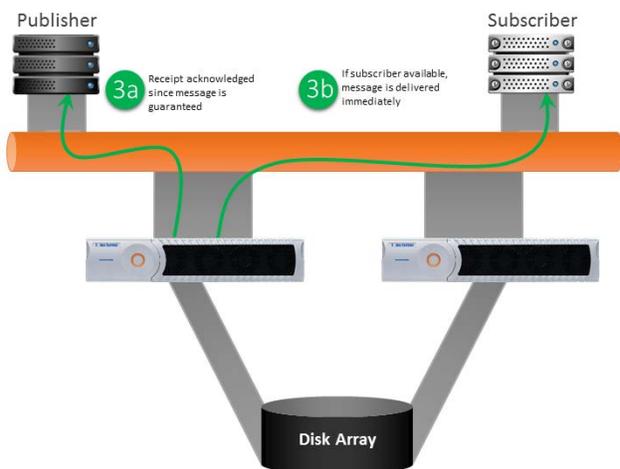


Acknowledgement of Receipt and Delivery

The router acknowledges receipt of the message back to the publisher once the message is securely stored, as the burden of guaranteeing the message now lives with the message router. If the subscriber is online and able to receive the flow of inbound messages, the message is immediately forwarded. When all of the consumers that need to receive a given message acknowledge receipt, the message router removes its copy of the message from memory.

Handling Slow / Disconnected Consumers

In the case of slow consumers, the message router pushes large blocks of data off to disk so the high performing RAM on the message router remains available for consumers that are online and keeping up with message flow. This way slow consumers never impact fast consumers, not even while slow consumers are catching up thanks to the message router's high capacity.



High Throughput, Low Latency

The only way software-based guaranteed messaging solutions can truly guarantee delivery, even through power failures, is to persist each message to disk. Disk writes are slow and add lots of latency – even for solid state disks. In fact, for latency-sensitive applications developers often accept the risk associated with turning off guaranteed delivery in their messaging system in order to achieve performance objectives. Their applications then have to take on the burden of assuring that no data is ever lost, which increases complexity and development costs and reduces application performance.

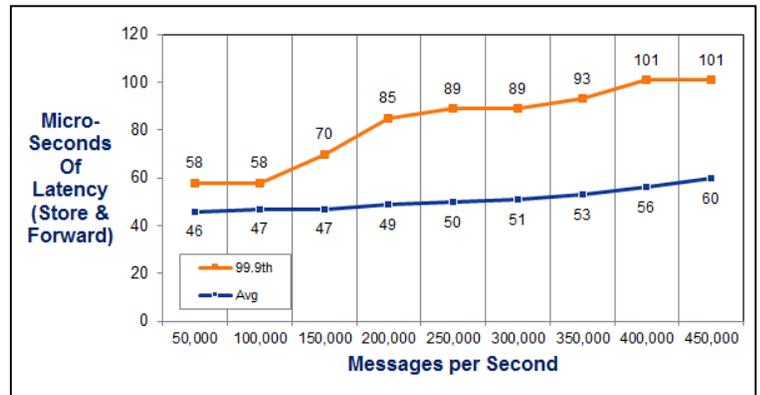
Solace's message routers store messages in high-speed on-board RAM as part of the data path unless recipients become unavailable to receive messages, at which time messages are batched in large blocks and stored to disk. Solace's patented techniques ensure that no messages, in memory or on disk, are ever lost, even in the event of a full power failure.

Replacing message-by-message persistence to disk with Solace's hardware-based failsafe delivery enables Solace's message routers to route over 450,000 messages per second, much higher than the few thousand that can be achieved in software. This eliminates the need to choose between fast and failsafe – for the first time, applications can have fully failsafe queued messaging at rates sufficient for most applications.

Solace supports two guaranteed messaging forwarding modes (cut-through and store and forward) so architects can choose the approach that best meets the functional and performance goals of each situation.

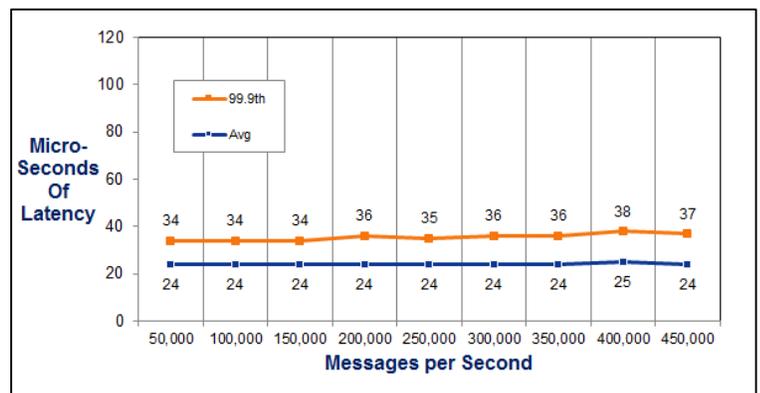
Store and Forward Performance

Solace can guarantee the failsafe delivery of messages using a store-and-forward approach that offers many sophisticated queuing features and low, consistent latency as shown by the graph to the right.



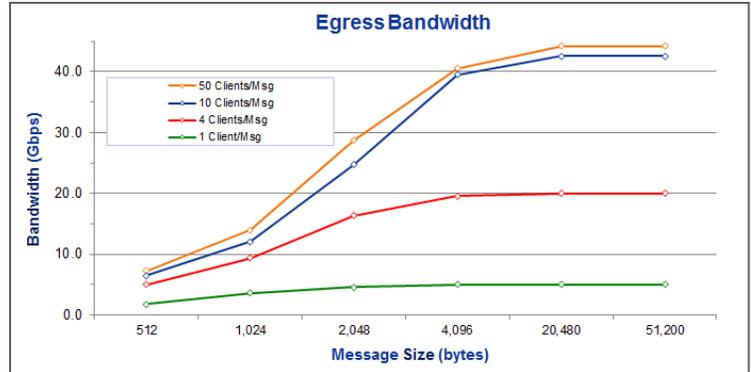
Cut-Through Performance

There are some scenarios, especially in capital markets trading applications such as smart order routing and crossing engines, where the value of the lowest possible and most consistent latency trumps the need for sophisticated functionality. For these situations, Solace offers the option of cut-through mode which enables guaranteed delivery with average latency of 25 microseconds and 99.9th percentile latency of 37 microseconds at 450,000 messages per second.



Fan-Out Performance

Solace message routers also offer extremely high performance when fanning out guaranteed messages to multiple consumers. This chart shows how egress bandwidth increases when messages of varying sizes are delivered to as many as 50 clients each – reaching an output bandwidth of 40Gbps.



Bulk Rate Performance

For applications where the latency requirement isn't so strict but overall message delivery rate is, we tested the message router for maximum message rate without fanout at various message payload sizes. These types of message rates represent the opportunity to achieve significant consolidation of messaging infrastructure and thus provide a more cost effective, simpler and robust infrastructure.

Bulk Message Rate	Message Size (bytes)	Message Rate (msgs/sec)	User Payload Bandwidth (Gbps)
	512	450,000	1.8
	1,024	434,000	3.6
	2,048	280,000	4.6
	4,096	152,500	5.0
	20,480	30,500	5.0

Steady Performance through Subscriber Disconnect and Recovery

SolaceMQ's deterministic high message rates and low latency solve another key problem with guaranteed messaging solutions – recovery from failure and the impact that can have on throughput rates. Consider an application that receives 1,000 messages per second. If it goes offline for just five minutes, 300,000 messages will need to be stored while it is offline and delivered when it comes back online. This is in addition to the 1,000 messages per second that continue to arrive while the application is recovering the old messages. In a software system limited to 3,000-5,000 messages per second for all applications using that broker, it can take hours to “catch up” and deliver those 300,000 messages, and service to other applications can be slowed or even interrupted during that time. Because Solace's guaranteed messaging solutions can deliver 450,000 messages per second, and has been designed specifically to provide robustness under these conditions, large backlogs can be metered into the system easily with no significant performance impact on existing applications.

To measure how well our message router isolates client applications from disturbances caused by other clients, we emulated a significant failure scenario—loss of connectivity and then restoring connectivity to more than a third of receivers (36%) — and measured what happened.

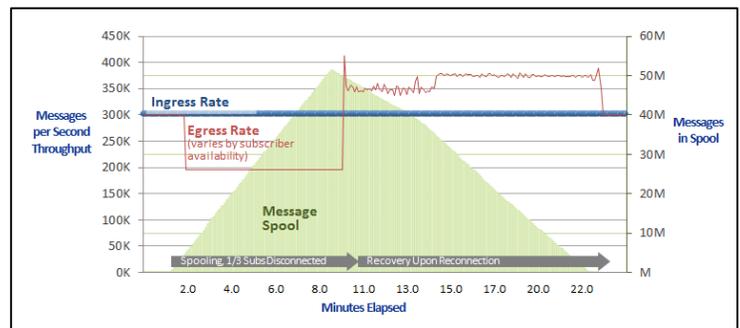
We started with 20 senders and 20 receivers at an aggregate target rate of 300,000 messages per second into and out of the message router. After two minutes, we disabled the Ethernet interfaces to 7 consumer processes and ran until 50 million stored messages had accumulated. Then we re-enabled the interfaces to all 7 failed consumers at the same time and continued to run the test until the backlog had been emptied. Then we ran for another five minutes. These tests were performed in a fully redundant environment using an active and standby message router to demonstrate performance in a typical High Availability deployment.

No effect on senders

Throughout the test the ingress message rate of the router (blue line on the graph) remained just above the target rate of 300,000 messages per second, showing that no senders were affected by the failure of the receivers.

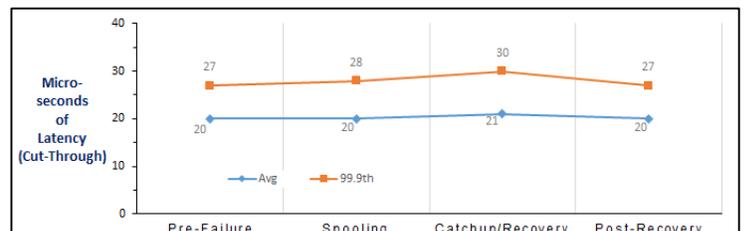
Throughput through dis- and re-connection

As far as the rate at which messages were received by receivers, there are four distinct patterns as shown by the black line on the graph: Prior to the disconnect, during the disconnection, during the recovery period, and after recovery. The egress message rate of the router is initially equal to the ingress rate at 300,000 messages per second. When the network disconnection occurs, 7 of the 20 receivers lose connectivity to the router and an equivalent 1/3rd drop in the egress rate was observed. When the connection is re-established, the egress rate peaks at 400,000 before settling in at 350,000 to 375,000 messages per second as the queued messages are released, and returns to 300,000 messages per second once all consumers have consumed their backlog and are back in real-time.



Latency impact on undisrupted subscribers

Next, we looked at the latency experienced by receivers that remained online throughout the test. Prior to the failure, performance was similar to previous steady-state tests. During the disconnection, while the router was storing messages on behalf of disconnected receivers, mean latency to undisrupted receivers was unchanged. During recovery, while the router was transmitting at a rate far in excess of the steady-state rate, mean latency rose slightly. Once the recovery period finished, the system returned to performance levels of the pre-failure period.



Virtualization

Each Solace message router can be partitioned into many “virtual messaging platforms.” This virtualization capability, in conjunction with Solace’s high capacity, means each one can replace many JMS or MQ brokers. Virtualization also lets administrators easily give applications and departments their own secure messaging environment. Messages never cross between environments, and administrators can configure the resource limits of each partition. You can even use this functionality to run development and test environments in one box.

This ability to horizontally scale within a compact footprint and share equipment can dramatically reduce TCO and time to market for new applications, and improve platform availability. To learn about Solace’s TCO advantage, visit <http://solacesystems.com/tco>



Other Advantages

Fault Tolerance and Immediate Failover for High Availability

When configured as fully redundant pairs for active-standby failover, the state of each message is synchronized across routers, so if the active router fails the standby system immediately recovers any in-process messages. To show how Solace guarantees the deliverability and sequence of messages in such a situation and provides for consistent fast failover, we ran a test in which we pulled the power plug on the primary router while under load. In this test we use 50 senders and 50 receivers at a rate of 100,000 messages per second. We then created an offline consumer that attracted all of the 100,000 messages per second into its queue until the message router had accumulated 350M messages and over 700GB of total payload storage. We then pulled the power plug out of the primary router and a failover occurred within 16 seconds. We then examined the results and found no data loss and that all messages were delivered in order. In the event of a catastrophic failure, Solace’s message routers have internally stored power to write all in-memory messages to an on-board flash disk. When power is restored, all messaging picks up where it left off. To see a demo, visit our YouTube channel.

JMS

Solace provides application access via the JMS API in addition to Solace APIs in Java, C# and C, and supports features including message selectors, time-to-live, dead message queue, queue browsers, dynamic and temporary queues. To learn more about the features and performance of our JMS solution, visit <http://solacesystems.com/jms>.

Unified Platform

Solace’s solution can be deployed as an extension of a Solace-based messaging environment and accessed by the same Solace APIs used for reliable messaging. By avoiding the need to deploy an entirely new software implementation to support persistent messaging, this keeps overall system complexity to a minimum, which reduces total cost of ownership.

Management & Monitoring

Rich information is provided on a per-connection and per-queue basis to provide insight into the operation of the system. In addition, various thresholds can be configured and alerts sent via SYSLOG and messaging when those thresholds are exceeded.

Security

Application authentication is provided via LDAP, Radius or internal database, along with access controls on a per topic and per-queue basis.