

Enabling Real-Time Sharing and Synchronization over the WAN

Solace message routers have been optimized to very efficiently distribute large amounts of data over wide area networks, enabling truly game-changing performance by eliminating many of the constraints imposed by today's solutions.

A quick summary of Solace's unique advantages:

- **WAN Optimized TCP:** Solace supports connection-specific TCP semantics (including a sliding TCP window) that greatly enhance behavior over high latency WAN links.
- **Hardware Compression:** Solace can automatically compress and decompress messages on a per-client basis entirely in hardware, reducing message sizes on average by 80%.
- **Parallel Traffic:** Solace enables the simultaneous delivery of many messages and acknowledgements over multiple TCP connections between Solace message routers.
- **Edge Fanout:** Solace message routers can automatically identify instances where many copies of a single message are required on the other side of a WAN link, then send just one copy over the link and fan it out on the other side.

This paper is designed to introduce application developers, network architects and IT decision makers to Solace's unique hardware-based approach to sharing and synchronizing data over long-distance network links.



The Challenge with WAN Data Distribution

Most messaging middleware products treat TCP connections as a service that the network layer provides, and don't adapt their behavior based on the different characteristics of LAN and WAN links. That means they handle messages the same way whether the network is providing 10 Gbps of bandwidth and microsecond latencies or 10 Mbps of expensive bandwidth with hundreds of milliseconds of latency.




In many cases, messaging applications connected over the WAN need to deliver data with a guarantee of receipt, and receive messages in the same order in which they were sent. Most products achieve guaranteed, sequential delivery by waiting for acknowledgement that each message has been received before sending the next message. This restricts throughput because the high latency of long-distance WAN links equates to very high round trip time (RTT). Worse yet, half of each RTT isn't being spent on the delivery of information, but rather the return of acknowledgement.

For example, if a given link's RTT is 100 milliseconds, and messages are always sent serially, your maximum message rate is 10 messages per second. Even if you double your bandwidth or cut message sizes in half, RTT will always be a gating factor that restricts message rate.

Fortunately, most commercial guaranteed messaging protocols (e.g. JMS and MQ) are smarter than that – when there are many messages waiting to be sent, they'll batch messages (preserving sequence) and send them as one trip across the network. This allows their effective throughput to exceed strictly serial limits.

Unfortunately, such optimization still leaves the majority of available bandwidth unused, because each batched send still waits for delivery and confirmations to complete before sending the next batch.

Customer testing has indicated that a connection on typical commercial JMS products use no more than 5-10% of the available bandwidth on a high latency 20-50 Mbps link, even if the data publish rates are much higher.



Solace's WAN-Aware Messaging

Solace's message routers have been optimized to be extremely efficient not just in LAN environments, but over the WAN. Here's a summary of Solace's unique advantages:

- **TCP Optimizations:** Solace's message routers allow connection-specific TCP semantics so high latency WAN applications can behave differently than high-speed LAN links. Over WAN links, Solace streams data through a sliding TCP window which allows messages to be continuously streamed, with delivery confirmations being streamed in the other direction. This alone overcomes the main throughput problem most commercial messaging products face, and allows much higher bandwidth utilization, often an order of magnitude higher.
- **Hardware Compression:** Solace message routers also perform compression in hardware, effectively reducing the size of messages on the wire by 80% on average. This allows 3-5 times more information to flow over the WAN link than the amount of data it could carry without compression. In many cases, applications can effectively exchange more than the data rate of the WAN link, extracting maximum value from the WAN investment.
- **Multiple Network Links:** By establishing multiple TCP connections between Solace message routers, traffic can be parallelized, which further increases throughput. Solace takes care of balancing traffic between the links so that messages are still sent and received in order while overall WAN traffic is balanced and throughput is optimized.




The combination of optimizing TCP, adding compression, and using multiple links dramatically improves throughput of messaging over a bandwidth constrained WAN link.

Optimizing High-Fanout Distribution

Solace further optimizes WAN distribution for applications requiring fanout to large numbers of recipients.

If a message has many subscribers across the WAN link, Solace's message routers send it just once over the WAN link, and fan it out at the remote location.



For example, consider database updates from Singapore that need to be shared with 20 back end systems in San Francisco, or a market data feed that originates in London and needs to be streamed to 100 traders in Tokyo. In these examples, the data rate savings can be an additional 20 and 100 times, respectively, in terms of WAN bandwidth consumed.

Solace's Advantage

The performance behavior and requirements of each WAN use case are unique. Depending on your budget, you may choose a 10 Mbps link vs a 100 Mbps link (or more). Between two nearby cities, your link may only introduce 30 milliseconds of round-trip latency, while a transoceanic link could add 300 milliseconds. Some applications need to exchange 100 byte messages while others sharing tens of kilobytes at a time. All of these factors, and many more, impact the behavior of your applications.

Regardless, Solace's WAN-optimized solution offers superior speed and resilience. Based on the gains we've seen real production customers experience, you can expect between 15 and 50 times more throughput than other commercial JMS or MQ products.

Distributed Management and Monitoring


Management and monitoring is especially important in WAN scenarios because the queuing of messages is so much more common. In the LAN, receiving applications don't usually fall behind the message stream unless the application is running slowly or disconnected. On the WAN, many more variables can cause messages to back up, such as inconsistent behavior of the WAN link and temporary volume spikes that exceed its capacity. Solace provides unique connection-level visibility into each endpoint's behavior and queues so operations teams have much broader and more granular visibility than software-based products give them.

Use Cases

The limited WAN throughput of JMS and MQ software products causes application designers and companies to think small when it comes to sharing and syncing data globally. They design business rules and build systems that assume their rate constraints and high latencies are unavoidable limitations of distributing information across datacenters and business units. As a result they design systems that share less data than they'd like to, share it less frequently than they'd like to share it, or write complex applications that have to work around these limits with batching techniques and horizontal scaling.

What kinds of applications would be possible if the messaging infrastructure could automatically scale, often beyond the theoretical maximum of the WAN link? And if you then buy more bandwidth, you actually get more throughput with no changes to the application – the messaging layer takes care of it all for you.

Below are some examples of the ways real companies are using Solace's WAN distribution and synchronization capabilities.



Distributed Big Data

The world of mobile devices, rich internet applications (RIAs), distributed sensor networks and cloud computing clusters require a new breed of data capture and analytics infrastructure that can handle the increasing volume and velocity of data.

Big data applications need to capture and consolidate data from remote locations for analysis by technologies like Hadoop or complex event processing engines. Then the insights gleaned need to be delivered to people and systems for further real-time adjustments. The volumes of data involved in big data applications have far surpassed the technologies used during the prior generation of distributed applications.

Big data is becoming the status quo in all kinds of information-heavy industries. For example, internet companies analyzing click-streams and traffic patterns, telcos tracing call/data activity, transportation companies keeping track of millions of packages and passengers, financial institutions managing transactions and risk around the world, and government agencies serving and protecting the public.


The best way to get big data flowing is with high-performance middleware that takes care of message queuing and delivery at very high rates so publishing applications and sensors can send data without worrying about where it needs to go or how it needs to get there. This entails the establishment and management of topics and queues, dynamic routing rules, and intelligent handling of fault conditions (such as applications or network links being down or slow).

Solace provides high-performance messaging solutions that excel at the distribution of large volumes of data. By handling routing and delivery in a pure hardware datapath, Solace supports higher throughput than software-based data distribution technologies with lower latency. Since Solace's solution is delivered as an easy to use message router, it's less complicated and costly to deploy, scale and operate.

Distributed Database Synchronization

To make smart decisions and run their business effectively, companies need accurate and timely information made available across their enterprise. Inaccessible and inconsistent data leads to bad decisions, dissatisfied customers and operational errors. Unfortunately, with data being generated and stored by countless devices and databases around the world, it's not easy to keep everything up to date. Most data replication and synchronization technologies are designed for periodic bulk transfers, don't make efficient use of expensive WAN links, and don't scale well because of the need to process routing rules in software.

Solace's WAN-optimized high-performance messaging solution, coupled with change data capture (CDC) software, enables the real-time synchronization of data across any number of geographically dispersed systems. In such a scenario the CDC software identifies data changes and passes them to Solace message routers which forward them to other systems that need that data. Solace's solution can route hundreds of thousands of guaranteed messages per second to any number of recipients, queuing them until delivery is confirmed so throughput is not restricted by the need to receive acknowledgements upon successful receipt of each update. The solution also performs message transformation as a service on the network so each system can send and receive information in its native format.



The best way to get big data flowing is with high-performance middleware that takes care of message queuing and delivery at very high rates so publishing applications and sensors can send data without worrying about where it needs to go or how it needs to get there.


By handling routing and delivery in hardware, Solace offers higher performance than software-based data distribution technologies with lower complexity and TCO.

Enterprise Data Grid

Driven by increasing business demands and the availability of technologies like in-memory databases, change data capture software, big data storage systems, and complex event processing engines, organizations are turning to enterprise data grids to accelerate, optimize and scale their IT infrastructure.

This new architecture is helping enterprises better compete in the global marketplace by improving the performance of application services and giving the business an agile architecture. This allows them to quickly create new revenue generating services, monetize their information assets and improve customer interactions.

Managing big data scale transactional and event stream information is about more than just storing massive amounts of data. It requires the intelligent collection, filtration, sharing and exposure of all that information via enterprise apps connected by LANs, WANs and cloud/grid environments. The diagram to the right shows the architecture of an enterprise data grid and components of next-generation enterprise-wide infrastructure.



Central to any enterprise data grid architecture is the Enterprise Information Distribution Fabric, the layer responsible for efficiently moving, filtering and synchronizing information between data grid instances, analytics and data warehouses around the world, all in real time.

Solace's message routers meet this requirement with higher performance, lower TCO and less complexity than any other solution—commercial, open source or home grown.

Summary

Whether you're synchronizing databases or data grids, fanning out content to remote users, or fanning in data from something like a global sensor network, Solace redefines what is possible over the WAN. For many customers Solace allows them, for the first time, to deploy WAN-based applications with expectation that the infrastructure can keep up with real time data change rates without breaking the bank in bandwidth fees.

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