
Financial Services: Inter-Data Center Applications

In the financial services industry, where companies routinely make sizeable dollar investments to shave a few milliseconds of WAN latency to maximize profits, inter-data center WAN traffic growth is top-of-mind for decision makers. The inter-data center WAN is witnessing dramatic traffic growth due to increased usage of technologies such as storage replication and virtualization, presenting financial services companies with both performance and economic challenges. The need to provision data centers hundreds of miles apart in order to prevent risks associated with natural and man-made disasters is further straining WAN resources.

Top Concerns of Financial Services CIOs

- Stronger data protection
- Faster time to market
- Accelerated market data

Optimizing inter-data center communications for financial services companies has its own unique set of requirements, the most important of which is the support for near real-time transactional applications, especially with Active/Active (multi-site, high availability) architectures that require changes to be written concurrently to multiple data repositories. Banks and brokerages process millions of dollars' worth of information in seconds; a few minutes or hours of downtime can be devastating. Beyond simple downtime, with Active/Active configurations, writes to each data repository must be carried out before the next transaction can be processed. Mindful of the need for business continuity in the strictest sense, financial institutions are locating secondary data centers farther and farther away from their primary data centers. Increased distance means increased latency, leading to performance degradation (e.g. for database writes) to keep up with synchronization requirements.

Backup and other data protection operations are increasingly being run over the WAN, while organizations retain increasingly larger amounts of data to comply with regulations such as Sarbanes-Oxley, the Basel Accords, and various others that govern customer data privacy. In addition, financial institutions must comply with regulations outlined by the Federal Reserve, the Office of the Currency, and the Securities and Exchange Commission (SEC) in their white paper titled "Interagency Paper on Sound Practices to Strengthen the Resilience of the US Financial System," which reflects lessons learned from September 11, 2001.

Although widely deployed, traditional WAN optimization controllers were not intended to respond to the challenges presented by inter-data center communications, and were instead optimized for branch office to data center communications. There are enormous differences between the two deployment scenarios: Branch office traffic is comprised of a large number of low-speed connections created by individual users, whereas inter-data center traffic is typically comprised of a small number of very high-speed connections provisioned between storage arrays, clusters, etc. This "machine" to "machine" communication is aggressive, bursty, and latency-sensitive, requiring high-bandwidth WANs.

These factors create a need for next generation WAN optimization devices that are purpose-built to support inter-data center traffic, not branch to data center traffic. This new class of WAN optimization devices must not introduce more than 100 micro-seconds of latency to the application workflow and must process data at wire speeds.

Common Financial Services Applications for Next Generation WAN Optimization Devices

Application: Storage replication for real-time production applications

In this case, latency is the biggest enemy. For example, synchronous storage replication promises zero loss of data because a write operation is not considered to be complete without an acknowledgement handshake between the storage systems at each data center. In most cases, the subsequent write operation will not start until the receipt of the acknowledgement of the previous write operation. Any significant latency caused by a long WAN link or introduced by a traditional WAN optimization controller can have grave consequences on the feasibility of synchronous storage replication as part of daily operations.

Application: B2C Web Sites and ATM Networks

Many financial services companies rely on distributed Web applications to provide customer portals for managing their accounts, transferring funds, or executing securities trades. For example, a large US-based bank may manage up to \$300 billion in assets with 3,000 branch offices and 7,000 ATMs. This bank must provide an always-on full line of banking, brokerage, insurance, investment, mortgage and payment service products to consumers, businesses and institutions. People around the world depend on ATMs for their banking needs anytime, anywhere; the availability of ATM networks is crucial to

Even a brief failure in a bank's ATM network or customer portal can result in a loss in customers and revenue.

today's functioning society. In addition, customers rely on constant online access to checking and savings accounts to receive payments and pay bills. As shown by recent incidents even a brief failure in a bank's ATM network or customer portal can cause a loss of account holder confidence, which may lead to lost customers and lost revenue for the financial institution.

These systems present great challenges to IT departments as they must be kept current within milliseconds in order to maintain transactional integrity. This is typically accomplished by replicating and synchronizing data between storage arrays located at two remote locations. For example, a multi-national bank may have regional data centers serving local branch offices. The company-wide customer financial accounts database must be kept up to date despite concurrent transactions taking place across the world. A typical solution is to execute transactions locally at each regional data center and then replicate data stores between data centers. This is typically accomplished using a solution such as EMC's Symmetrix Remote Data Facility (SRDF), a replication solution for EMC's high-end SAN arrays. The bank would have multiple high performance storage systems, each capable of managing petabytes of data. Transactions will be executed against the closest regional data center, saved to the local SAN, and replicated across WAN links asynchronously.

Using disk to disk replication across fast WAN links instead of disk to tape backup has helped many financial services companies cut their recovery time from 16+ hours to less than one hour. However, this places significant and ever-increasing demands on WAN links. If latency and bandwidth aren't managed carefully, asynchronous replication can fall behind and data back ends end up out of sync. The higher the latency and the lower the utilized bandwidth, the more likely this inconsistency or detrimental break is to occur.

Application: Backup for Business Continuity and Disaster Recovery

Backups, whether they are disk to disk or disk to tape, are critical components of business continuity and disaster recovery strategies. These processes protect physical and virtual application servers, databases and file storage – the core systems provided by IT departments. Data is growing at an alarming rate, as is the need to push more data between data centers for the purposes of business continuity and disaster recovery. Every organization wants to back up data faster and minimize backup windows. Typically, backups occur in the evening when there is less production traffic traversing the WAN. When backup volumes increase dramatically, backup windows grow in tandem. With operational hours cutting into backup time budgets, there is increasing pressure to either purchase more bandwidth (to remain within the backup windows), to move backup processes into the daytime, or worse, to simply backup less data.

72% of IT departments at Fortune 500 banks / financial institutions have discarded critical data from backups in order to reduce backup windows.

Criteria for Selecting Next Generation WAN Optimization

Next generation WAN optimization solutions should manage and accelerate critical applications requiring high performance and throughput between data centers. These solutions should be able to keep the WAN link fully utilized while applying reduction features to reduce the actual data footprint on the WAN.

Core Selection Criteria:

- **Acceleration:** The device must be able to accelerate up to 10Gbps of traffic.
- **Data Reduction:** The device must be able to sustain inline data reduction rates at levels of 80-90%.
- **High Throughput:** The device must support up to 1Gbps per connection and guarantee high throughput for critical workflows traversing long distance (100ms+ latency) WAN links.
- **Quality of Service:** The device must support granular QoS policies to ensure tight SLAs for critical application traffic.

Next generation WAN optimization devices need to be built to sustain multiple 10 Gbps links with extremely low latency. WAN traffic rates needs to be reduced, and network overhead needs to be dealt with more efficiently. This requires a very fast hardware-based reduction engine that enables the WAN optimization device to maintain the highest levels of data reduction at multi-gigabit speeds, while guaranteeing port to port latencies in the tens of microseconds.

TCP windows and inter-flow behavior needs to be managed very aggressively with dynamically adjusted window sizes in order to curtail the effects of WAN latency and packet loss. And finally, the WAN optimization controller must be built around a very high speed traffic management sub-system that can classify and prioritize traffic at multi-gigabit line rates, so that competing traffic can be managed as per business priorities.

About Sarrel Group

Sarrel Group is a technology product assessment, custom publishing, and IT consulting firm with offices in New York City and San Francisco. Sarrel Group helps technology companies develop sales and marketing programs based on lab tested validation of their product's competitive advantages. Sarrel Group has been instrumental in the development, evaluation, and implementation of thousands of technology solutions.



For more information, please visit www.sarrelgroup.com or call 866-MSARREL.