

Growing Margin Lending while Reducing Intraday Trading Risk in Real-time

Mitigating risk, enhancing customer experience, and achieving unlimited growth potential with mainframe offloading



Results achieved with Aerospike

- Migrated end-of-day batch system to real-time, customer facing solution significantly improving customer satisfaction while saving \$10,000 per day with mainframe offloading
- Consolidated 150 server caching layer workload with 12 Aerospike servers resulting in 90% TCO reduction
- Perform 250 million transactions and 2 million price updates per day
- 400% faster throughput with 300% larger object store
- Data consistency mainframe-to-cache
- APIs from browser-to-server

The Company

One of the world's top three global brokerages, known for its robust online and mobile trading platforms and commission-free trades, offers full-service investment management, online trading and banking services to clients around the world with total client assets north of \$3.5 trillion, was growing over 10 percent annually and had also recently acquired a competitor.

Background

The firm's IT group was at a crossroads. The combination of an RDBMS and a RAM-based cache fronting a traditional mainframe database was unable to consistently and reliably support ever-growing workloads during trading hours. This reality had become a major impediment to the company's strategic goal of releasing a steady flow of new and updated applications to its mobile customer base. The business mandate was to seamlessly service more than 10 million customers and process more than 250 million transactions daily, while positioning itself to eventually process 1 billion transactions per day, was at risk.

The company's conventional RDBMS infrastructure and cache capabilities couldn't consistently deliver the performance, scalability, and flexibility needed to meet these requirements. In its prior configuration, the caching solution quickly became overloaded, rendering it unable to scale during high-volume trading days or handle new applications without serious latency issues and system outages. System restarts were taking as much as an hour to complete, and given its reliance on RAM, which isn't persistent, data loss was a serious problem.

Overall, the firm depends on delivering a better customer experience in a highly competitive marketplace. As such, they were interested in moving high impact, customer-facing solutions to real time from existing end-of-day batch processing. In addition, they saw potential opportunities to increase revenue through more focused risk management to extend appropriate margin limits while better monitoring client portfolio risk.

"We've not had issues where customers complained about something not processed correctly on the backend system. It's a true reflection of the capabilities that Aerospike provides."

VICE PRESIDENT, TECHNOLOGY
TOP 3 GLOBAL BROKERAGE FIRM

Business requirements

To scale without barriers, provide superlative customer experience, and continue to introduce new, engaging mobile applications, this very large brokerage had to rethink its infrastructure. Continuing with a relational database and cache-based solution as the intraday system of record would require scaling from 150 servers to more than 1,000 servers. This was not a practical strategy in terms of time, labor, and operating cost. More important for the business, reliance on nightly batch processing from the intraday system to the master DB2 (book of record) was expensive, cumbersome, not customer friendly, and did not solve the problem of data inconsistencies between stored and active data.

A new solution had to address:

- The company's decision to continue to leverage its legacy mainframe database (DB2), which was the compliant system of record for more than 10 million customer accounts.
- The requirement to process 250 million transactions and 2 million updates a day, and the ability to update stock prices or show balances on 300 million positions in near real time.
- The ability to create enough compute capacity to eliminate data inconsistencies.
- The elimination of frequent system crashes due to overloading the RAM-based cache and the subsequent restarts, which frequently took around an hour.
- The mandate for a cost-effective solution that would address expectations for 1,000% data growth as it executed on its mobile strategy.

In working with the Vice President of Technology, the business requirements were clear. According to the vice president, "When we are dealing with our customers' money, activities and transactions, we need to make sure that it's never an override, an exception, or any wrong data getting posted," they stated.

Beating the Competition

The company's IT department explored a number of options and quickly realized that Aerospike was the only solution that could provide the required performance due to its unique hybrid architecture. Rather than use expensive yet unreliable RAM, the firm preferred Aerospike's ability to leverage solid state drive (SSD) flash technology. The firm realized flash pushes down costs and increases reliability because it has much greater storage density than RAM (thus requiring fewer servers) and immediate data persistence.

The brokerage evaluated multiple competitors including MongoDB, Gigaspaces, and Gemfire, which were eliminated from consideration for a variety of reasons:

- MongoDB couldn't provide low read latency at a high write load. It was also not flash-optimized so it could not significantly reduce server count.
- Gigaspaces is an application-centric cache solution and guarantees varied based on the database layer choice. Its architecture tied applications to the distribution layer, making it impossible to efficiently access data from multiple clients.
- Gemfire needed at least 4x more nodes and could not deliver predictable performance with low latency at high throughput.

Delivering better business results with Aerospike

Even with the increased load during the COVID-19 crisis, the brokerage has been able to meet and exceed customer expectations and there has been no disruption in Aerospike's stability, latency and performance. The vice president of technology even refers to Aerospike as "the secret sauce" when it comes to low latency for the organization.

Following are some ways that Aerospike is helping this Top 3 Global Brokerage firm deliver better business results:

Making real-time customer portfolio data available to drive margin loans

The company's previous legacy system was handling both the reads and writes, which was a problem since a typical day saw 80 percent of the interactions as "reads" (customers checking on portfolio performance, reviewing market gains or losses) and 20 percent as "writes"(customer trades), the Vice President of Technology indicated. That meant that high read ratio was causing a lot of stress on the legacy system. The solution: setting up Aerospike clusters and building an Aerospike integration layer, which talks to the systems of record whenever the legacy applications are writing to the database. Now, all the data in the system of record is available in real time in a consumable format. "We are now able to perform over 13 billion queries per day with 99.9 percent requests returned in sub-milliseconds," he says. "Based on our ability to keep all intraday event-based data available – with strong consistency and high throughput – in the application layer, we are now able to offer margin loans with appropriate risk management and compliance controls for over 14 million accounts in real time – the interest of which is a major source of income for (our company)."

"Based on our ability to keep all intraday event-based data available in the application layer (with Aerospike), we are now able to offer margin loans for over 14 million accounts in real time – the interest of which is a major source of income."

**Vice President, Technology
Top 3 Global Brokerage Firm**

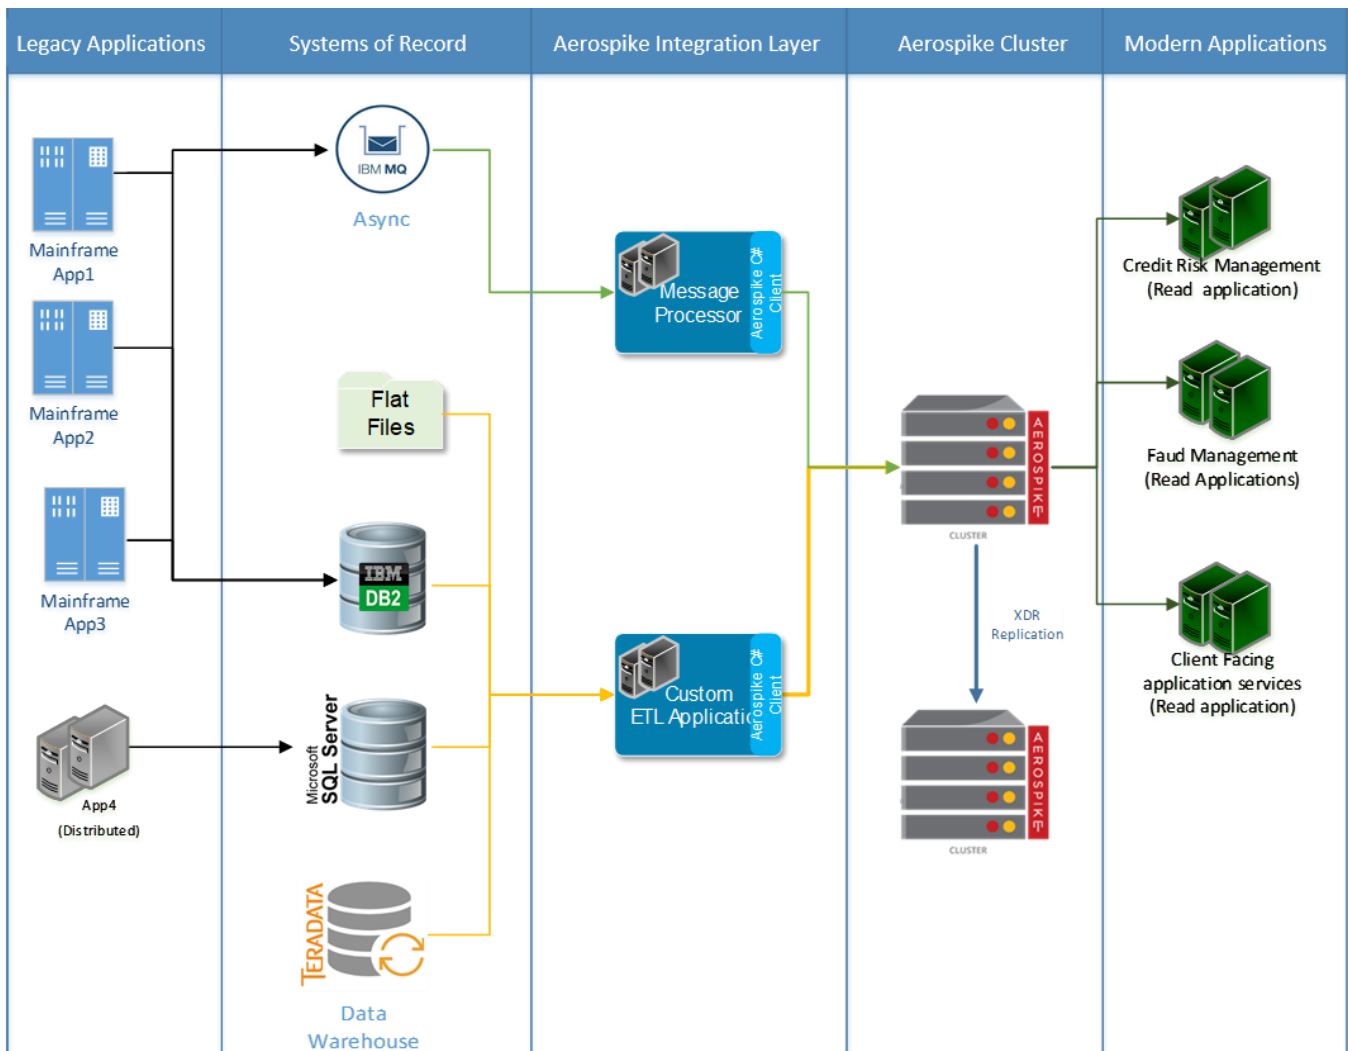


Figure 1 – Aerospike as Cache for Read Applications

Implementing a new intraday operational data store

Several years ago, the organization wanted to move the more complex "writes" – trading and orders – onto the modern layers. "So now, when publishing an event to a legacy system asynchronously, we are writing to a queue process that will update the source system, a system of record," the VP of Technology says. "It is a reverse paradigm of the read application." Now, Aerospike layers get data at the beginning of the day. The site application reads the data, then writes it into Aerospike within the same transaction. It then posts a message into the messaging queue.

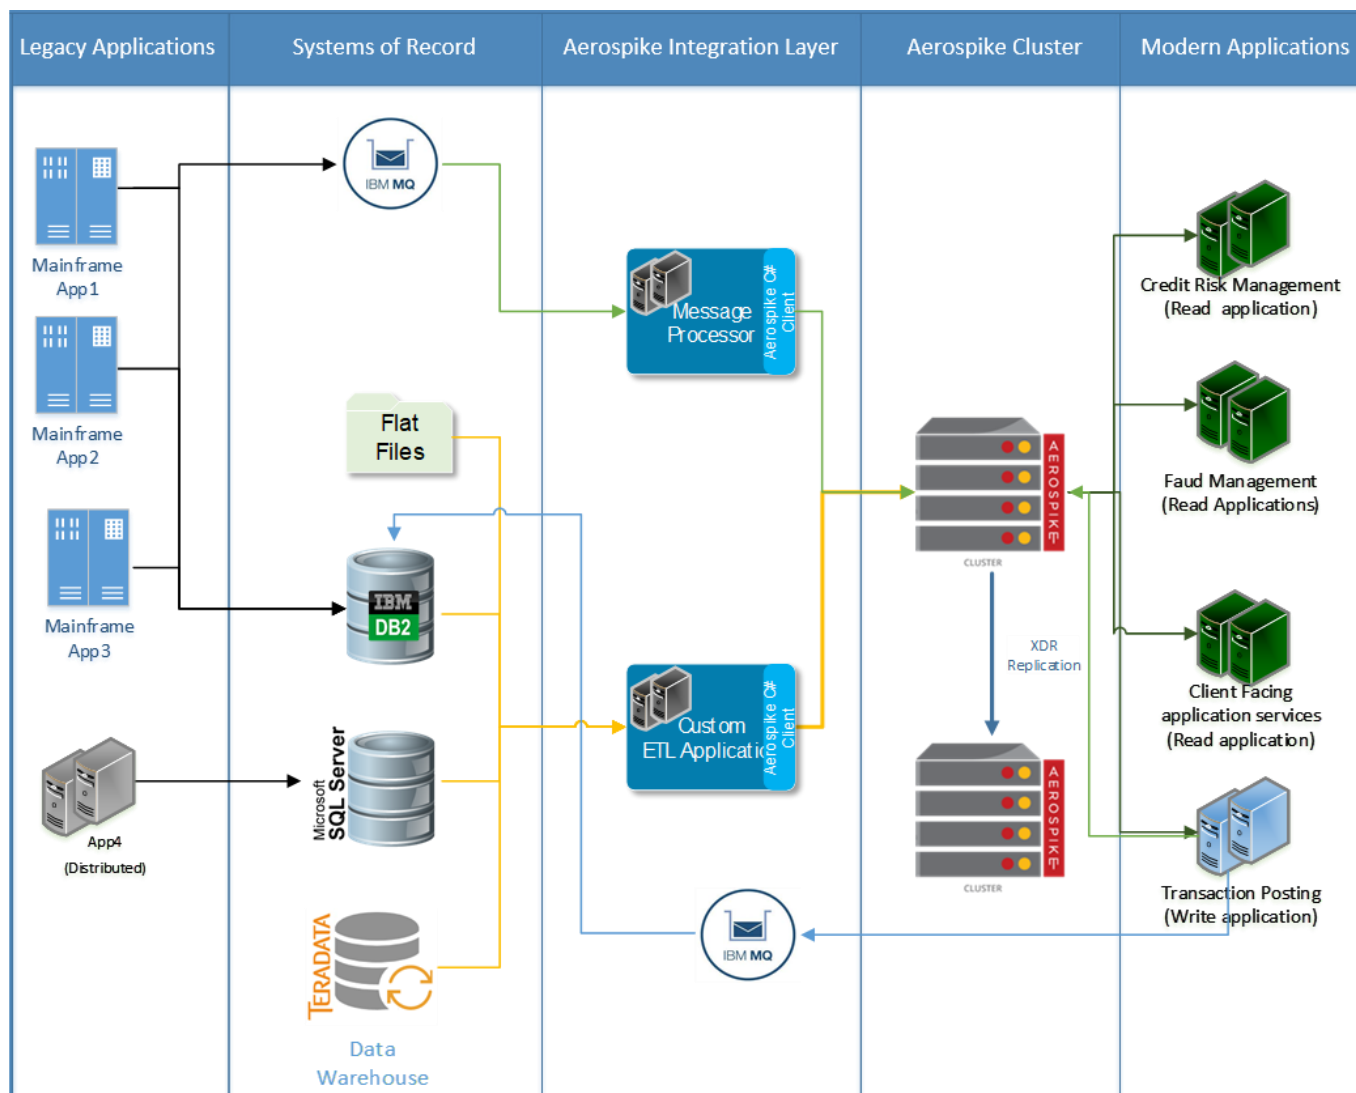


Figure 2 – Aerospike as an intraday operational data store

Using Aerospike as a true operational data store during intraday and as its system of record, the Global Brokerage is now benefiting even when a subsystem is not available. Stated the vice president, "If there is an outage or any other issues, the customer experience is not impacted. Messages will queue up and when they become available they will be written, keeping the system of record as a source system with our mainframe DB2 because this data is leveraged by lots of other applications."

Aerospike serves as the intraday system of record, while DB2 still functions as the compliance related 'Book of Record' and synchronizes with Aerospike at the end of each trading day. Because Aerospike runs natively in flash, which is significantly faster and more dense than disc, it makes dramatically more efficient use of the hardware resources. As a result, the company reduced its hardware footprint from 150 RAM cache servers to a 10-server Aerospike cluster, while simultaneously improving performance.

Empowering new capabilities by providing easy access to vast amounts of data

“We continue to get all the benefits of low latency, with all needed historical data available for us to access it in real time, so that we keep up with our SLAs with the customer channels.”

**Vice President, Technology
Top 3 Global Brokerage Firm**

Since Aerospike acts as the system of record, this means there is real-time access to all data. This makes it easier to scale and manage and gives the Global Brokerage firm the opportunity to leverage backup and recovery capabilities. For example, one recent use case involved the brokerage firm’s need to retain large amounts of historical data of over 500 terabytes of readily accessible, available data. The company worked with Aerospike to develop an advanced hardware strategy to allow them to store up to six terabytes in one disk, “coming up with an impressively small footprint,” the Vice President of Technology said. “And we continue to get all the benefits of low latency, with all needed historical data available for us to access it in real time, so that we keep up with our SLAs with the customer channels.”

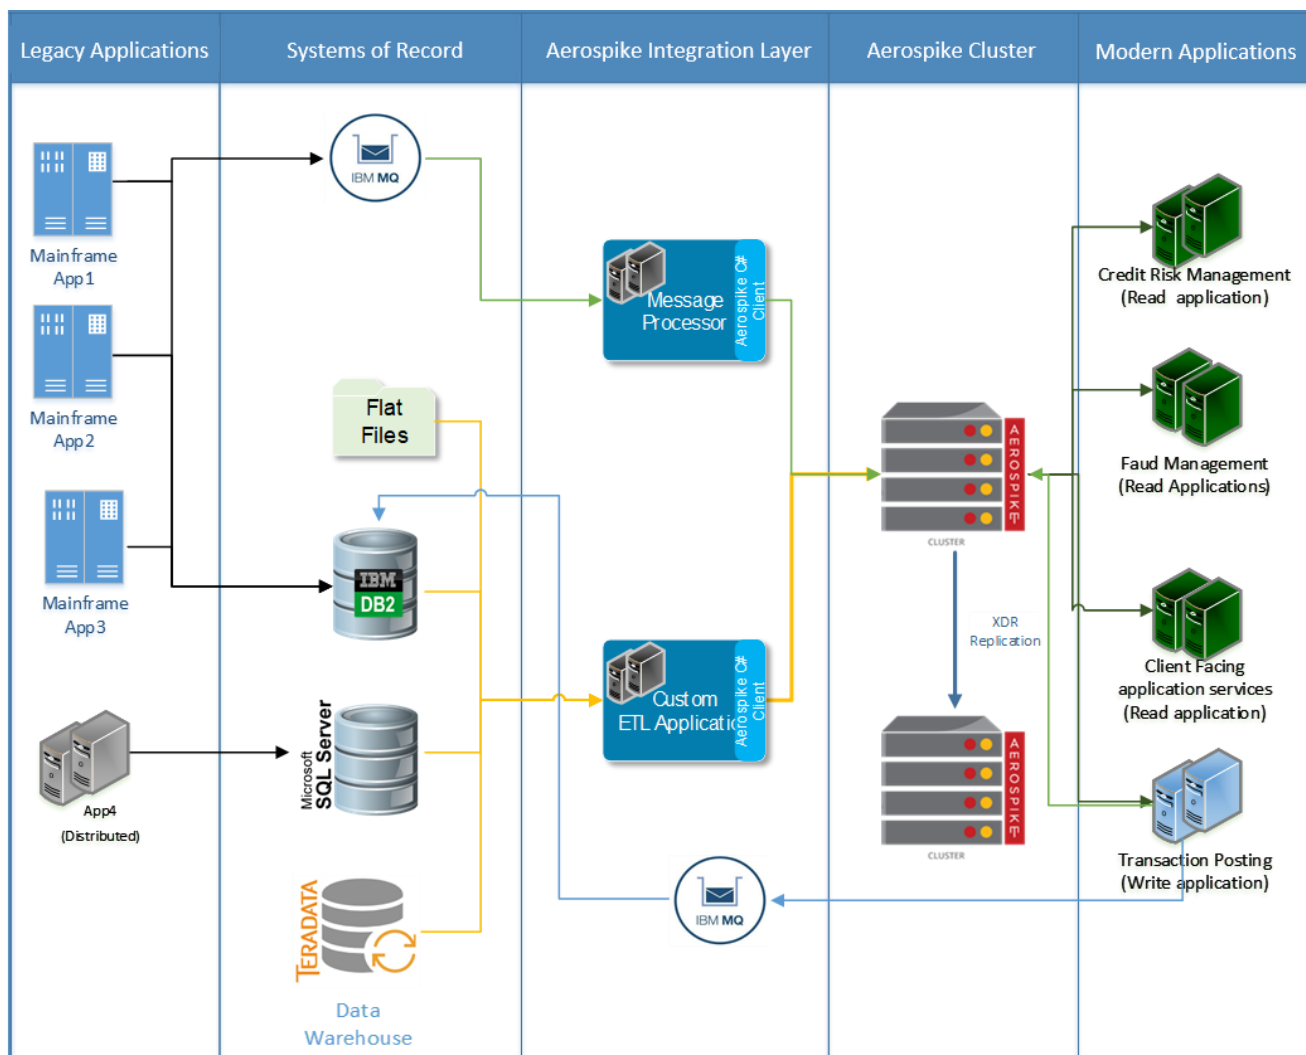


Figure 4 - New Applications that generate vast amounts of data - Aerospike as System of Record

The Vice President of Technology says that Global Brokerage firm also benefits from leveraging many of Aerospike's features, including multi-site clusters and XDR (cross data center replication) capabilities as the company runs in multiple, distributed data centers that have multiple network cores. Since the firm is running no shared architecture in terms of deployment, that means that whatever happens in one data center or one network core should not impact anything on the other data centers.

"Leveraging those capabilities has gotten us to the point where we are able to support shared architecture at the same time, providing all the latency and availability requirements," he explains.

Going forward with Aerospike

Regarding Aerospike's reliability and accuracy over the many years deployed, the VP of Technology indicated, "We've not had issues where customers complained about something not processed correctly on the backend system. It's a true reflection of the capabilities that Aerospike provides."

"It has been an exciting ride at (our company) for the last seven years, especially as we acquired (one of our competitors)," the VP of Technology said. "Scalability is on top of mind right now across every organization and across every application. Whatever applications we migrated over (in the above three use cases), we're confident that it is only adding more hardware – this (Aerospike solution) will scale as much as needed."

Finally, as the firm continues its modernization journey, it will continue to rely on Aerospike to provide not only the most reliable and cost-effective solutions, but also a valued strategic partner.

About Aerospike

Aerospike is the global leader in next-generation, real-time NoSQL data solutions for any scale. Aerospike enterprises overcome seemingly impossible data bottlenecks to compete and win with a fraction of the infrastructure complexity and cost of legacy NoSQL databases. Aerospike's patented Hybrid Memory Architecture™ delivers an unbreakable competitive advantage by unlocking the full potential of modern hardware, delivering previously unimaginable value from vast amounts of data at the edge, to the core and in the cloud. Aerospike empowers customers to instantly fight fraud; dramatically increase shopping cart size; deploy global digital payment networks; and deliver instant, one-to-one personalization for millions of customers. Aerospike customers include Airtel, Banca d'Italia, Nielsen, PayPal, Snap, Verizon Media and Wayfair. The company is headquartered in Mountain View, Calif., with additional locations in London; Bengaluru, India; and Tel Aviv, Israel.