



Highlights

- Accelerate business-critical financial applications with IBM infrastructure
 - Achieve record-breaking performance¹ with the highly efficient eXtremeDB Financial Edition
 - Help lower costs and reduce complexity with IBM® Power Systems™
 - Leverage the ultra-low latency of IBM FlashSystem® to gain competitive advantage
-

Powering the financial industry

IBM FlashSystem and IBM Power Systems: Proven performance for financial systems

The financial services industry encompasses a wide range of businesses, from traditional hometown banking, through investment and wealth management, to global securities exchanges. As in almost every other industry, the digital revolution has brought great change to the financial sector. Perhaps the greatest change has been in the speed at which financial business now operates.

Banking systems are getting faster as they move from pure systems of record to systems of engagement, driven by new online and mobile applications. In the 21st century, equities sales occur in milliseconds, with ultra-high-speed algorithmic trading on the rise as tick data response times decline. Like never before, the performance of back-end IT infrastructure directly affects the ability of financial institutions to succeed.

IBM has provided technology solutions to the financial industry for more than 100 years. Today, 96 of the world's top 100 banks and nine out of 10 of the world's largest insurance companies run on IBM IT infrastructure.² As part of its on-going commitment to provide financial enterprises with the best IT solutions available, IBM has recently teamed with embedded database technology company McObject and the Securities Technology Analysis Center (STAC)³ to perform benchmark performance testing on a reference architecture specifically designed to accelerate financial applications that included McObject's *eXtremeDB* Financial Edition 7.0 hosted on an IBM Power Systems server with



IBM FlashSystem storage. In one standardized benchmark, which runs 21 different tests, the IBM infrastructure delivered the fastest and most consistent mean response times ever reported for all combinations of query type, data volume and concurrent users. In another test, the IBM infrastructure beat the performance of the best result previously published by a stunning 200 times.¹

The STAC testing confirms that deploying McObject's *eXtremeDB* in an environment of Power Systems and IBM FlashSystem can provide financial institutions with the application performance they need to capture and keep competitive advantage, increase profitability and move confidently into what IBM calls the "Cognitive Era."

Securities Technology Analysis Center

STAC provides technology research and testing tools based upon community-source standards. STAC facilitates the STAC Benchmark Council,⁴ an organization of leading financial institutions and technology vendors that specifies standard ways to assess technologies used in finance. The council is active in an expanding range of low-latency, big-data and big-compute solution stacks for financial data processing. STAC helps end-user firms relate the performance of new technologies to that of their existing systems by supplying them with standardized STAC Reports as well as standards-based STAC Test Harnesses for rapid execution of STAC Benchmarks.

High-performance solution test bed

The STAC-M3 test with McObject's *eXtremeDB* demonstrates that Power Systems servers and IBM FlashSystem all-flash storage provide an exceptional low-latency and low-jitter implementation of tick data warehousing and analysis. The extensive cache/memory hierarchy provided by IBM Power Systems S824L and the performance of the IBM FlashSystem 900 array rival even the most optimized systems based on other CPU architectures with large storage clustered arrays. The jitter advantages of this configuration and McObject's *eXtremeDB* demonstrate an enterprise-ready tick data warehouse never before available.

McObject

McObject provides database technology to meet the unique needs of both big-data systems and intelligent, connected devices. Founded by veterans of the database and real-time systems industries, the company has developed *eXtremeDB*, a small-footprint database system optimized for applications where high performance and low latency are critical. *eXtremeDB* supports both in-memory and persistent databases and can be used as an embedded or client/server database system.

eXtremeDB

eXtremeDB Financial Edition achieves record-breaking performance¹ due to its highly efficient architecture, hybrid row/columnar storage, database sharding, powerful indexing capabilities, extensive vector-based math functions and calculation pipelining for on-chip analytics. As a true ACID-compliant database management system for big data, *eXtremeDB* Financial Edition scales predictably and is both robust and highly reliable, characteristics that are critical for a wide variety of capital market applications where a rich feature set coupled with high performance, low latency and minimal jitter are important. *eXtremeDB* has been used for more than 15 years in high-performance applications in capital markets, defense and avionics, among others. There are hundreds of companies using the development software and more than 30 million copies deployed in applications around the world.

The volume of data flowing through today's automated capital markets is skyrocketing, and success for financial technology hinges on acting on this information instantly. However, alternatives today lack the performance required by low-latency applications, are too difficult to scale and/or are too expensive. *eXtremeDB* Financial Edition leverages a streamlined hybrid in-memory/persistent storage database design, multi-core optimization, maximum developer flexibility and high scalability to address key financial data management challenges to make the solution an extraordinarily fast database system with extreme flexibility for financial applications.

IBM Power Systems

To power the STAC benchmark testing with McObject's *eXtremeDB*, IBM chose the Power Systems S824L server with its advanced microprocessor technology. Power Systems server platforms deliver the performance and scaling capacity required by 21st century financial applications. With new innovations, Power Systems provide the ability to:

- Move data in and out of systems more quickly with twice the memory and input/output (I/O) expansion
- Achieve greater speed and efficiency for database, online transaction processing (OLTP), and other highly multi-threaded applications with processing memory supported by 50 percent more cores and twice the number of simultaneous threads per core

The IBM POWER8® microprocessor is the fourth generation of IBM POWER® technology announced in 2001 as IBM POWER4, the first multi-core server technology in the industry. IBM Research has contributed chip-manufacturing technology that delivers high reliability, high quality, high-density memory and high-bandwidth interconnects. The IBM-patented embedded DRAM (eDRAM) enables the POWER8 processor to increase cache density to double the per-core on-chip capacity of competitive processors.⁵ In addition, POWER8 servers can deliver up to 128 MB of L4 cache per socket, which is shared among the processor memory interfaces. This feature alone dramatically improves the performance of persistent database systems' cache as well as that of time-series databases such as *eXtremeDB* Financial Edition. POWER8 also provides simultaneous multi-threading (SMT) with up to eight threads per core. Additional scale-up POWER8 servers increase that number to 192 cores in a single operating system image of a nearly flat SMP model. Power Systems S824L is the ideal system for distributed scale-out tick data workloads.

Power Systems are ideal for consolidating multiple applications and infrastructure workloads in a virtualized environment such as those in which financial applications are most often deployed, bringing together business transaction processing with infrastructure for big data, analytics and OLTP solutions.

IBM FlashSystem

An IBM FlashSystem 900 all-flash array was chosen as the back-end storage target for the STAC benchmark tests because it provides extremely low response times from a single, easily-deployed platform. IBM FlashSystem 900 and its precursor models have been deployed in the most business-critical financial application environments for nearly a decade.



Figure 1. IBM FlashSystem 900

Powered by IBM FlashCore™ technology, IBM FlashSystem 900 delivers the extreme performance, enterprise reliability and operational efficiencies required by the financial industry and by *eXtremeDB* users. For storage environments equipped with robust storage management capabilities such as those provided by *eXtremeDB*, the ultra-low latency, market-leading performance and ease of implementation offered by IBM FlashSystem 900 make it an ideal choice for transforming business-critical applications into engines of business growth and competitive advantage.

IBM FlashSystem 900 is composed of up to 12 massively parallel IBM MicroLatency® modules that provide very high storage density with response times in the 100 microsecond range. IBM FlashSystem 900 can scale usable capacity from as low as two terabytes (TB) to as much as 57 TB in a single 2U array. It supports AES-256 encryption and comes with high-speed internal interfaces and full hot-swap and storage capacity scale-out capabilities, enabling organizations to achieve lower cost per capacity with the same enterprise reliability.

Benchmark test design and configuration

Leading global financial institutions developed the STAC-M3 Benchmark suite to represent a cross-section of real-world use cases. In the study conducted with IBM and McObject,

two STAC-M3 Benchmark suites were executed—Antuco and Kanaga. The Antuco version uses a relatively limited data set, while the Kanaga benchmark employs more scalable (varying number of concurrent requests) queries over larger data sets. As illustrated in Figure 2, these tests executed the STAC-M3 Benchmarks on a test bed stack that included McObject’s *eXtremeDB* Financial Edition 7.0 hosted on an IBM Power Systems S824L server with two 12-core POWER8 CPUs connected to a 12 x 5.7 TB IBM FlashSystem 900 storage array. Power Systems S824L executed 72 threads on a 24-core

Test bed configuration

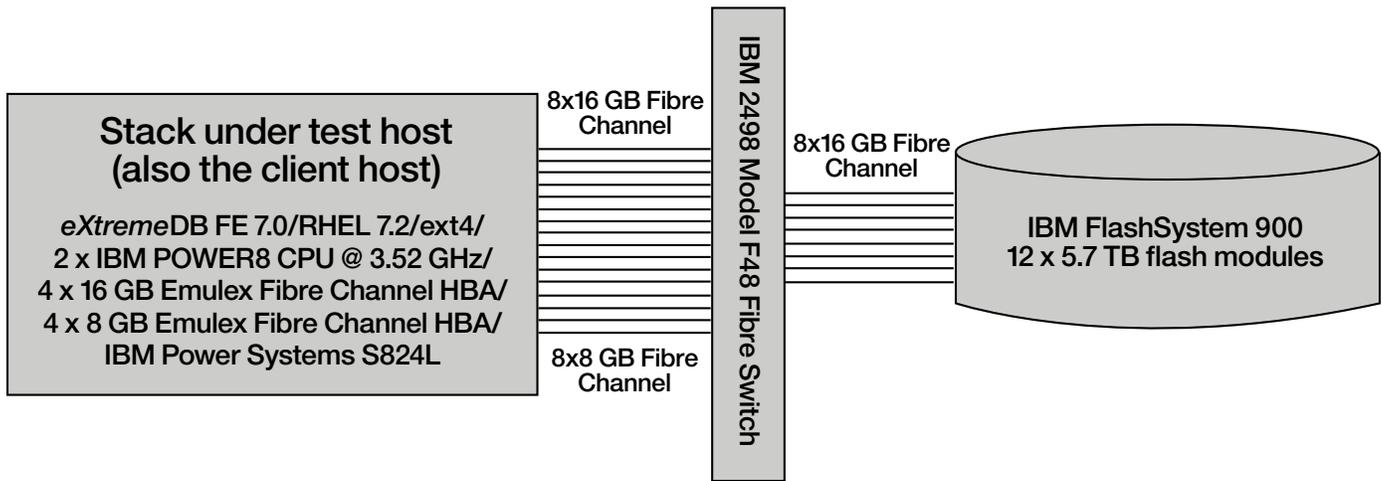


Figure 2. Test bed configuration used for STAC M3 testing

(2-socket) system to achieve record-setting performance.¹ This stack also included a new revision of the *eXtremeDB* STAC Pack, which enhanced performance.⁶

For this test effort, McObject implemented the STAC-M3 clients and operations using the STAC-M3 Benchmark specifications. IBM supplied the lab and support, server and storage hardware. McObject and IBM configured and optimized the full stack under test. McObject sponsored the audit. Finally, STAC conducted the STAC-M3 Benchmark audit, which included validating the database, inspecting any source-code revisions to the STAC Pack, validating the operation results, executing the tests and documenting the results.

Overview of the STAC-M3 Benchmarks

As STAC puts it, analyzing time-series data such as tick-by-tick trade and quote histories is crucial to many trading functions, from algorithm development to risk management. But the domination of liquid markets by automated trading—especially high-frequency trading—has made such analysis both more urgent and more challenging than ever before. As trading robots try to outwit each other on a microsecond scale, they dish out quotes and trades in ever more impressive volumes. This places a premium on technology that can store and analyze that activity efficiently. For example, the faster algorithm developers can back-test and discard a haystack of unprofitable ideas, the faster they will find the needle of a winning algorithm, leaving more time to exploit it in the market.

The STAC Benchmark Council has developed STAC-M3 Benchmarks in order to provide a common basis for quantifying the extent to which emerging hardware and software innovations improve the performance of tick storage, retrieval and analysis. STAC-M3 tests the ability of a solution stack to perform a variety of operations on a large store of market data. The STAC-M3 Working Group designed these test specifications to enable useful comparisons of entire solution stacks, as well as comparisons of specific stack layers while holding other layers constant.

Stack under test setup

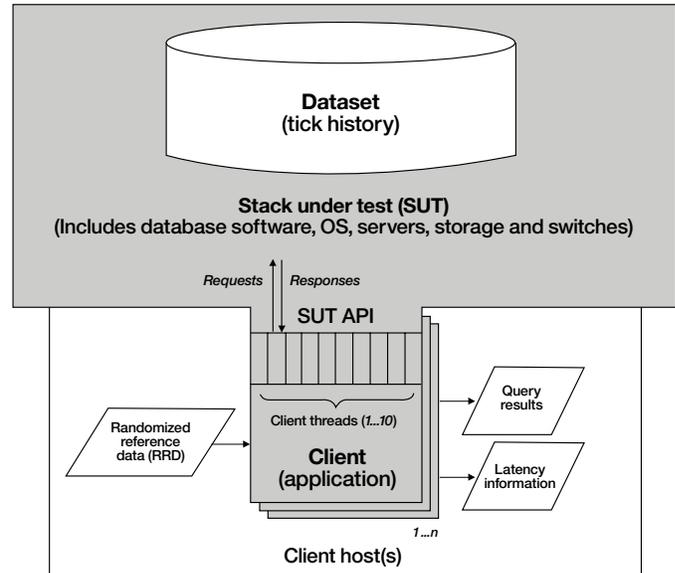


Figure 3. Standard STAC M3 stack under test (SUT) setup

As shown in Figure 3, the test setup for STAC-M3 consists of the “stack under test” (SUT) and client applications. No restrictions are placed on the architecture of the SUT or clients. Threads within the clients take in randomized reference data (RRD) such as dates and symbols, submit requests for the required operations, receive responses, and store the timings and results from these queries. Vendor-supplied code for the operations and latency calculations are subjected to a combination of source-code inspection and empirical validation.

STAC-M3 Antuco and Kanaga tests

The test setup for the STAC-M3 Kanaga suite is identical to that for the STAC-M3 Antuco suite. The STAC-M3 Kanaga data set is an extension of the STAC-M3 Antuco data set, which is a synthetic store of trades and quotes modeled on one year of New York Stock Exchange trade and quote data for US equities. Since it represents only one year of one data set, the Antuco database is considerably smaller than the full data sets in use at customer sites. This was a deliberate choice by the STAC-M3 Working Group to minimize the cost of running the baseline benchmarks while still yielding valuable results.

The optional Kanaga suite of STAC-M3 consists of test sequences that extend Antuco benchmarks across larger quantities of data in order to measure the volume scalability of a database stack, as well as a large number of users. The ability of a tick database solution to handle increasing volumes of historical data is important to today's trading organizations. Market data volumes continue to grow quickly, sometimes in step-function increments. Engineering a solution that delivers consistent, high performance across ever larger data sets can be a challenge. STAC-M3 Kanaga provides some insight into how well a given tick database solution scales.

The Kanaga suite runs tests with up to 100 concurrent simulated user requests or more and requires a minimum storage of more than 30 TB. It allows additional, future years to be added to the database, using the same data structures, where each year's size is a multiple of the previous year's size. Using a fairly aggressive rule of thumb that assumes total market data volumes double roughly every 18 months, Kanaga sets the annual scaling factor at 1.6. In addition to scaling forward in time, the specifications also scale the data backward several years in order to test how the technology stack handles older, smaller data sets that may be on slower archive media.

Extraordinary results

McObject's *eXtremeDB* running on a POWER8 and IBM FlashSystem 900 infrastructure yielded impressive results when executed against both the STACM3 Antuco and Kanaga test suites.¹ Against the Kanaga suite in particular, *eXtremeDB* powered by IBM FlashSystem and POWER8 achieved:

- 21 times to 212 times the performance of the previous best published result for the market snap benchmarks
- 21 times the performance of the previous best published result for "year high bid" in the smallest year of the data set
- 8 to 10 times the performance of the previous best published result for the 100-user volume-weighted average bid benchmarks
- 5 to 8 times the performance of the previous best published result for the N-year high-bid benchmarks

Against the Antuco suite, the IBM and *eXtremeDB* test bed delivered:

- 1.93 times the performance of the previously published best result for 10T.THEOPL.TIME
- Faster mean response times than the previously published best results for four of the 17 operations
- 1.3 times the performance of the previously published best result for 1T.STATS-UI.TIME
- Compared to all publicly reported two-socket systems, the system set records in four of the 17 benchmarks; twelve of the remaining records are still held by previous *eXtremeDB*/POWER8 systems (SUT IDs XTR141023, XTR151016) series of tests
- More consistent response times (lower standard deviation) than the previous best results for four of the 17 operations

Accelerating the financial industry

In the financial industry, competitive advantage is measured in milliseconds and customer experience. Whether racing to complete equities trades faster than competitors or ingesting, processing and analyzing a flood of online customer transactions, financial institutions demand the highest possible performance from their IT infrastructures.

STAC benchmark tests confirm that an IT solution architecture composed of McObject's popular *eXtremeDB* database powered by IBM FlashSystem all-flash storage and Power Systems processors can cost-effectively provide financial enterprises with outstanding application performance.

For more information

To learn more about IBM FlashSystem and IBM Power Systems, please contact your IBM representative or IBM Business Partner, or visit: ibm.com/flash

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition. For more information, visit: ibm.com/financing



© Copyright IBM Corporation 2016

IBM Systems
Route 100
Somers, NY 10589

Produced in the United States of America
August 2016

IBM, the IBM logo, ibm.com, Power Systems, POWER8, IBM FlashSystem, IBM FlashCore, and MicroLatency are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

¹ Based on STAC Benchmark testing conducted in 2016.

“STAC Report: eXtremeDB & IBM at scale under STAC-M3,” May 9, 2016. <https://stacresearch.com/news/2016/05/09/XTR160413>

² Janet L. Sun, “Don’t believe the myth about the mainframe, Part 1,” SHARE, Inc., May 7, 2013. <http://www.share.org/p/bl/et/blogid=2&blogaid=234>

³ STAC and all STAC names are trademarks or registered trademarks of the Securities Technology Analysis Center LLC.

⁴ For more information on the STAC Benchmark Council, visit www.STACresearch.com/council

⁵ Based on IBM internal lab testing.

⁶ The code used to execute these benchmarks is available from the STAC Vault for inspection and use by qualified STAC Benchmark Council members.



Please Recycle