

FINANCE INDUSTRY DATA ON A NoSQL SOLUTION

MARKLOGIC WHITE PAPER • JANUARY 2016

Moving the financial firm's data from legacy mainframe to Enterprise NoSQL on Intel®-based solutions



THE OPPORTUNITY FOR TRANSFORMATION

The data needs of the finance industry are changing. The system of record is being asked to provide more thorough search and query capabilities for regulatory compliance and e-discovery. All the while, the intense requirements for security, reliability, and capacity for high-volume transactions remain.

The mainframe architecture has served many of the archiving, transaction, and system of record needs of the finance industry for an impressively long time.

However, as businesses look to the new horizons of data applications, utilizing the legacy mainframe is a costly and time-consuming hurdle. In addition, the increasing demands of regulation require more sophisticated records, which can extend beyond the capabilities of an older database.

Although maintaining the current installed system may seem to be the most affordable decision for the immediate financial cycle, analysis of the insidious increases in maintenance costs along with lost opportunities builds a strong case for migrating sooner rather than later.

The legacy mainframe is a cost center and slows data innovation. Organizations are losing ground every day the mainframe remains at the center of the enterprise. The trend toward reduced complexity, along with platform consolidation and environment rationalization, makes a compelling case for considering Intel® Xeon® processor-based solutions.

A TALE OF TWO TIMELINES: BITEMPORAL DATA FOR THE FINANCE INDUSTRY

Traditional databases are unitemporal, meaning that they store data along one timeline. If you fail to create a method of storing further data such as *when you changed the data* then critical information can be lost – information that can save energy in compliance situations and provide analytic insight.

A bitemporal database can record timestamps for events along two dimensions of time: valid time and system time. Valid time tracks when an event occurred

in the real world. System time (sometimes called “transaction time”) tracks when the event is recorded to the database.

For the finance industry, this can be a significant savings of effort and compliance risk. With a more realistic record of the time of events, organizations can avoid entanglements with regulatory or legal discovery.

As anyone responsible for record keeping in the financial industry knows, improving the integrity of those records for regulatory and e-discovery is a critical issue. Along with the increasing cost of regular compliance practices, the financial institution faces an increasing risk of fines. In the United States, banks paid a record-breaking \$65 billion in regulatory fines in 2014¹.

BEFORE BITEMPORAL	AFTER BITEMPORAL
What do we think the trader's position was, and what information do we think was available to the trader around the time when the trade was executed?	What was the trader's exact position when the trade was executed, and what exact reference data was available at the time the trade was executed?
What were our customer's credit ratings last year?	What were our customer's credit ratings last year, as we knew them last quarter?
What was our market exposure when trade was made at 11:00am?	What was our market exposure when that trade was made at 11:00am, as we knew it at 11:30am?
What was the company's profit when we gave guidance?	What did we think the company's profit was when we gave guidance?

¹ James Sterngold, “For Banks, 2014 Was a Year of Big Penalties”, Dec. 30, 2014 <http://www.wsj.com/articles/no-more-regulatory-nice-guy-for-banks-1419957394>

DATA IS THE FINANCE INDUSTRY'S COMPETITIVE EDGE

Big data is the topic of conversation across all industries, and for good reason. As far back as 2011, a McKinsey & Company report found that the leaders in big data usage consistently outperformed the rest of their industry in both revenue and earnings². In a 2015 survey of executives from large global companies, 90 percent had medium to high investment in big data. Two-thirds reported that big data and analytics have had significant, measurable impact on revenues³.

This push for big data solutions has led to more and more business-unit requests to conceive of information in a new way. This new way of looking at the business can stretch the resources of the mainframe. In many cases, the task is entirely beyond the capabilities of the existing system of record.

The strategies that drive the modern finance industry require additional transaction and data capabilities. The new standard expects flexible, inclusive, and distributed data.

THE GROWING COST OF DOING NOTHING

Future opportunity can be hard to quantify, but current costs are a serious and measurable issue. While there is a short-term cost to adopting a new technology capable of flexible and creative data implementation, the status quo will be increasingly more expensive. The maintenance of a legacy mainframe is becoming costlier by the year.



As an organization's legacy system ages and widespread use of mainframe technology fades further into history, it will cost more and more to maintain the same level of service.

² "Big Data, The Next Frontier for Innovation, Competition, and Productivity," McKinsey Global Institute, 2011.

As the skillsets to maintain the equipment and codebase sunset, programmers and administrators are more difficult to find and increasingly expensive. Following a major data breach, the United States Office of Personnel Management released its 2014 Strategic Information Technology Plan. Along with concerns about security, the report prioritizes migrating away from the legacy mainframe systems because cost of support is expected to rise 10–15 percent annually⁴. And this doesn't even take into consideration opportunity costs of not being able to use mainframe data in new big data solutions, such as real time data analytics. And this doesn't even take into consideration opportunity costs of not being able to use mainframe data in new big data solutions, such as real time data analytics.

THE ALTERNATIVE: DISTRIBUTED DATA CENTERS

While the administration and operation costs of legacy mainframes, including maintenance and new applications, are getting higher, the capabilities and return on investment (ROI) of Intel-based alternatives increase. The Intel-based data center solution ecosystem is being revolutionized by innovation in management, energy efficiency, and performance. Modern Intel-based data centers are able to run more services per staff hour, megawatt, and clock cycle than ever before.

INTEGRATED DATACENTER TECHNOLOGY

Intel storage, network, and compute technologies are increasing the productivity and performance of data centers.

- **Scalable Processors Designed for Data Centers**
For example, the Xeon E5-2600 v3 product family, which is designed for the workloads of a data center, includes numerous enhancements that provide performance increases up to three times that of the previous generation⁵, world-class energy efficiency, and enhanced security.

³ Press, Gil, "6 Observations from a New Survey on the State of Big Data Analytics," Forbes.com, September 4, 2015. Found at <http://www.forbes.com/sites/gilpress/2015/09/04/6-observations-from-a-new-survey-on-the-state-of-big-data-analytics/>.

⁴ Strategic Information Technology Plan, U.S. Office of Personnel Management, 2014 Ftics/.

⁵ As of September 8, 2014. New configuration: Hewlett-Packard Company HP ProLiant ML350 Gen9 platform with two Intel Xeon Processor E5-2699 v3, Oracle Java Standard Edition 8 update 11, 190,674 SPECjbb2013-MultiJVM max-jOPS, 47,139 SPECjbb2013-MultiJVM critical-jOPS. Baseline: Cisco Systems Cisco UCS C240 M3 platform with two Intel Xeon Processor E5-2697 v2, Oracle Java Standard Edition 7 update 45, 63,079 SPECjbb2013-MultiJVM max-jOPS, 23,797 SPECjbb2013-MultiJVM critical-jOPS <http://www.spec.org/jbb2013/results/res2014q1/jbb2013-20140121-00050.html>.

- **Intelligent and Efficient Storage** – The Intel SSD DC P3608 Series solid state drives (SSD) deliver high-performance and low latency. Non-Volatile Memory Express (NVMe) and eight lanes of PCIe 3.0 eliminate bottlenecks in high performance computing workflows, accelerate databases, and provide business insights through real time analytics. The unique dual controller architecture scales efficiently with multiple core Intel Xeon processors.
- **Integrated Network Systems** – With 1 billion Intel Ethernet ports shipped and more than 30 years of innovation, Intel offers a comprehensive portfolio that includes controllers, adapters, IP, and switch silicon. Since the inception of Ethernet, Intel has continued to help drive Ethernet technology by providing technological advances and active leadership in standards committees and industry associations.

NOSQL: SKIPPING ETL AND SAYING YES

For decades, the art of data has been designing database schemas and the craft of extract, transform, and load (ETL) processes. Standardizing addresses, financial currencies, and other fields is traditionally a key part of any new data integration or consolidation project on a legacy mainframe system of record. Mainframe data often needs to be copied into a different application-centric database or data warehouse, requiring an ETL process.

ETL projects need to be vetted for accuracy. The content and process need to be kept up-to-date. ETL can be resource intensive, both for staff and compute. For the business units, ETL is cost, delay, and constraints.

The transform portion in particular is resource intensive and complex. In the transforming of data, some aspects of the information can be lost or compromised. For instance, if a building height is calculated using different parameters in different databases and those need to be stored together, the variances might be lost.

When an organization adopts a NoSQL approach to data management, the database imports content in its native state. The ETL process is replaced by the significantly less involved “extract and load.” Sometimes termed “schema-agnostic,” Enterprise NoSQL document databases do not require extensive design, planning, or architecture before introducing a new data source.

ENABLING EMERGING BIG DATA APPLICATIONS

With modern processing and data tools, organizations are discovering new value in their existing data sets. They create dynamic correlations by combining dissimilar data sources. Enterprise NoSQL is a compelling enabler of future data applications. An opportunity that requires the consumption of JSON, XML, unstructured files, text, RDF triples, geospatial, or any binary can be implemented rapidly without the pre-work of ETL and schema design.

A million heterogeneous documents can be corralled to answer a question like “where are sales stumbling in Europe?” No matter the format each business unit uses, data and documents from a variety of sources can be combined to maximize service availability.

MarkLogic / Enterprise NoSQL Database Platform

POWERFUL	AGILE	TRUSTED
<ul style="list-style-type: none"> Native JSON Store Native XML Store Native RDF Triple Store Geospatial Support Full-text Search Flexible Indexes Bitemporal Real-time Alerting Semantic Inference Tiered Storage Server-side JavaScript Fully Transactional 	<ul style="list-style-type: none"> Scalable and Elastic Cloud Ready (AWS) Hadoop and HDFS REST API SQL Support Multi-OS Support Schema Agnostic Samplestack MarkLogic Content Pump XA Transactions Ad-hoc Queries Index Across Data Types 	<ul style="list-style-type: none"> Performance at scale LDAP and Kerberos Security Security Certifications Monitoring and Management Configuration Management 24/7 Engineering Support ACID Transactions Flexible Replication Customizable Backup Customizable Failover Point-in-time Recovery Atomic Forests

Enterprise NoSQL is a trusted platform that cost-effectively responds to modern business expectations.

MISSION-CRITICAL FINANCE APPS ON NoSQL

Designed for durability and scalability, Enterprise NoSQL technology from MarkLogic supports massively scaled applications around the world. MarkLogic® Enterprise NoSQL is the platform for organizations in need of a system of record capable of high volume and high flexibility. Many companies are already realizing the benefits of MarkLogic Enterprise NoSQL.

DISTRIBUTED AND FLEXIBLE DATA CENTER ARCHITECTURE

When dealing with branch offices, new service offerings, and changing business climates, organizations need to be able to affordably grow and transform IT infrastructure. With a mainframe, one can add processing and storage, but adding secondary locations is not so easy. Architecting for bursts and peaks means essentially purchasing compute for the maximum need.

TOP 5 BANK

32 million live deals	1,600 requests per second
-----------------------	---------------------------

BBC STREAMING 2012 OLYMPIC COVERAGE

55 million global browsers	2.8 petabytes
----------------------------	---------------

HEALTHCARE.GOV

5,000 highly complex transactions a second	11+ million Americans with health coverage
--	--

CONDÉ NAST

7+ million rich media assets	<1 second search results
------------------------------	--------------------------

Organizations adopting Intel® Xeon® processor-based data centers can manage growth and opportunity using modern flexible computing techniques. Virtualization of servers and infrastructure allows for an unprecedented ability to maximize resource use.

Bursts in demand and new initiatives mean that an IT organization needs to be able to cost-effectively increase compute capacity, perhaps for only a brief period of time. Applications developed on Intel-based solutions have options. Along with reallocating resources within the data center, hybrid cloud solutions can extend the application to a platform as a service (PaaS) or an infrastructure as a service (IaaS) vendor. MarkLogic is architected to handle elasticity and new distribution deployments without downtime.

ENTERPRISE NoSQL + INTEL-BASED SOLUTIONS

- Are ready for new integrations
- Are designed for distributed workloads
- Scale up and down elastically
- Lower management costs
- Lower development costs
- Lower TCO costs

MIGRATING THE SYSTEM OF RECORD FROM A LEGACY MAINFRAME TO NOSQL

Achieving the promise of a new technology usually means a migration process. Migrating a system of record, the heartbeat of a company, is no small task and not one to be taken lightly. As the market moves away from mainframe solutions, actionable guidance and help is available to the organization that embarks on a migration path.

Intel offers best practices for the migration planning and execution in its paper “Migration from UNIX/RISC and Mainframe to Intel-Based Solutions.”

HSBC RE-HOSTS MODERN MAINFRAME APPLICATIONS

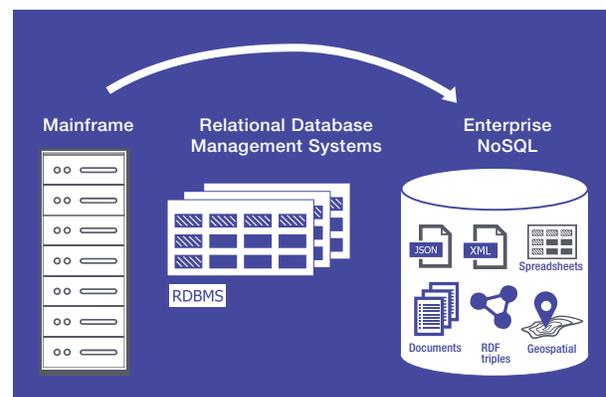
HSBC encountered performance and reliability issues for a key mainframe after deploying a new suite of loan applications. Migrating the applications to Intel Xeon processor-based servers delivered:

- A better experience for customers and branch office personnel due to improved performance and higher uptime
- Lower costs and greater headroom, with 70 percent lower monthly service charges and a 2,000 MIPS reduction in mainframe workloads
- Improved business flexibility with an infrastructure that is easier to scale and adapt

Read the full case study for more information about the HSBC methodology and migration experience at <http://www.intel.com/./data-center-xeon-hsbc-whitepaper.pdf>

For more guidance on migration to an Intel architecture data center, visit the Intel [IT Center for Server Migration](#)

LEAPFROG RDBMS: ENTERPRISE NoSQL IS A MORE COMPELLING DESTINATION



Migrating to Enterprise NoSQL provides greater capabilities and a less disruptive migration path than RDBMS.

The legacy mainframe applications that are still running today have most likely survived despite previous

challengers. Throughout the years, IT departments have evaluated migrations to relational database management systems (RDBMS). Many times, they chose to leave the mainframe in place.

Since the 1990s, relational databases have become the expected database solution for most data projects. Still, there is a huge base of applications still running on thousands of mainframe systems around the world. They were not migrated for very good reasons: migration is costly, disruptive, and risky.

Enterprise NoSQL, along with offering extended capabilities, has a safer, less intensive migration path and offers greater data agility to ensure the data is future-proofed for the next 40 or more years. With a more flexible information storage, less transformation means less time and less risk.

IT CAN BE DONE; IT HAS BEEN DONE: COPYBOOKS TO XML

While data is stored and managed differently in a modern NoSQL database, the XML-friendly structure allows for some logical migration paths. Ken Krupa, MarkLogic CTO, utilized the COBOL from his early career to [document a process](#) that will reliably ingest COBOL copybooks into the MarkLogic Enterprise NoSQL database.

CASE STUDY: POSITIONING FOR THE NEW COMPLIANCE LANDSCAPE

Today, a global broker-dealer is faced with a complex challenge; they must capture and retain sophisticated data from multiple sources and be able to efficiently recall that information upon demand for regulators.

In 2012 a top-tier global investment bank set out to consolidate its compliance data from multiple source databases around the globe. The goal was a permanent archive of relevant business records necessary for investigations and regulatory responses. They selected an XML-based solution built upon MarkLogic Enterprise NoSQL and scaled it to support up to a petabyte of data.

Collecting data from sources that included legacy mainframes, aged Sybase servers, and modern trading applications, they were able to integrate 40 applications in the first 12 months of implementation and provide unprecedented response to legal requests.

ONGOING FEATURES:

The bank's new compliance system is designed to expand with new features going forward, including:

- **Tiered data storage**
- **Operational trade repository**
- **Reference data management**
- **Regulatory and legal compliance**
- **Customer analysis and insights**
- **Cybersecurity and fraud prevention**
- **Pre-trade decision support**
- **Information distribution**

MEASURING SUCCESS

For most business units, a reliable system of record is the expected baseline. Surveyed IT departments have found that business units can be the source of the most resistance to a migration. This is a reflection of the importance of reliability. Job one is to preserve or enhance the customer experience in line of business applications.

For IT departments and executives, the financial bottom line is an obvious gauge of success. Reducing the cost of support and of developing new applications can be a path to realizing return on investment for the migration efforts.

From compliance to customer service to real-time market analysis, Enterprise NoSQL on Intel-based solutions is an enabling technology for future financial industry applications: those currently planned and those not yet imagined. For the future, an organization can realize the value of an integrated data repository that provides access to all aspects of the organization's data profile. An Enterprise NoSQL solution is designed to quickly integrate new data sources as they appear, no matter the format or source.

ABOUT MARKLOGIC

For more than a decade, MarkLogic has delivered a powerful, agile, and trusted Enterprise NoSQL database platform that enables organizations to turn all data into valuable and actionable information. Organizations around the world rely on MarkLogic's enterprise-grade technology to power the new generation of information applications. MarkLogic is headquartered in Silicon Valley and has offices throughout the U.S., Europe, Asia, and Australia. For more information, please visit www.marklogic.com.

ABOUT INTEL

Intel (NASDAQ: INTC) is a world leader in computing innovation. The company designs and builds the essential technologies that serve as the foundation for the world's computing devices. As a leader in corporate responsibility and sustainability, Intel also manufactures the world's first commercially available "conflict-free" microprocessors. Additional information about Intel is available at newsroom.intel.com and blogs.intel.com, and information about Intel's conflict-free efforts is at conflictfree.intel.com.