

GAINING TRANSPARENCY WITH CHANGING REGULATIONS

WHITE PAPER • APRIL 2017 – V1

There is opportunity in the ongoing turbulence throughout the financial industry; however, regulatory and compliance challenges absorb resources and distract focus away from growth. A smart approach to data and innovative tools build resilience into a business looking to not just weather the storm but seize the advantage.



EXECUTIVE SUMMARY

The cost of compliance for financial services organizations is exploding: more regulations to address, more transparency to provide, faster reporting requirements, and rapid change is the only constant. But the cost of noncompliance can be even higher. Fines, litigation, and impact to company reputation are at stake for companies that can't keep up.

At the same time, money is on the move because of growing uncertainty across regional and global markets. This presents great opportunities for organizations flexible enough to maintain the velocity of both transactions and regulations. As great as the risks can be, the rewards for those who can keep up with the pace can be greater.

Meeting today's demands isn't easy, and it can't be done the old way. Data silos; slow and expensive extract, transform, and load (ETL) processes; and traditional database architectures don't give you the flexibility you need to meet compliance requirements and to take advantage of the opportunities in the marketplace today.

MarkLogic's unique information management capabilities running on Intel's industry-leading open and scalable architecture provide the financial services sector with a sustainable route out of the complex morass of regulatory investigations and reporting. Organizations need to break down traditional data barriers and enable a 360-degree view of data across the company. You need flexibility to meet tough compliance and security requirements while improving productivity and profitability.

THE CHALLENGES OF THE FINANCIAL MARKETS TODAY

The financial services industry is under constant pressure from financial regulation bodies in all regions of the globe. Institutions are subject to requirements, restrictions, and guidelines by a variety of regulators. In recent memory, organizations have had to accommodate regulatory changes from the Prudential Regulation Authority (PRA) and the Financial Conduct Authority (FCA) in the United Kingdom (overseeing the implementation of European reporting requirements such as FINREP and COREP); the Basel Accord; and the Sarbanes-Oxley Act of 2002 in the United States.

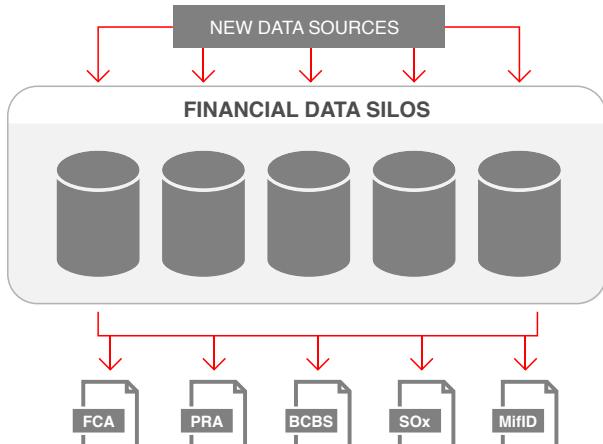


Figure 1 – Regulatory agencies and new data sources put stress on heterogeneous data silos

As part of these requirements, the financial institutions are obliged to report periodically—daily, monthly, quarterly, or annually—to the various regulators. These regulatory reports contain raw or summary data needed by the regulators to evaluate the safe and sound condition and operations of a bank or other financial institution or to determine compliance with any government act or other law, rule, or regulation.

This input is a vast amount of data from a broad variety of sources. It includes data about trades executed, reference, and counterparty data all sourced from trading and market data systems. The analysis of that data and the speed with which it needs to be manipulated require considerable and complex processing resources.

Furthermore, the banks are obliged to ingest new data sources and correlate them with existing data to provide even deeper levels of insight for the regulators. Not only do the banks have to report periodically to the regulators, they can be subject to ad hoc reporting demands for stress and scenario testing. To support the regulatory framework, the regulators can set policies for firms' stress testing requirements, set stress scenarios, and monitor the test results.

To meet the needs of today, financial institutions need a 360-degree view of data throughout the organization, need to increase transparency and agility, and need to decrease the time required to ingest and integrate new data as it flows into the organization.

BCBS 239 RISK DATA AGGREGATION AND REPORTING



Figure 2

As an example, the Basel Committee on Banking Supervision (BCBS) requires documentation and monitoring across a broad array of business and IT activities. Basel III institutes new requirements.

- I. Governance, data architecture, and IT infrastructure
- II. Risk data aggregation capabilities
 - Accuracy and integrity
 - Completeness
 - Timeliness
 - Adaptability
- III. Risk reporting practices
 - Accuracy
 - Comprehensiveness
 - Clarity and usefulness
 - Frequency
 - Distribution
- IV. Supervisory review, tools, and cooperation
 - Review
 - Remedial actions and supervisory measures
 - Home/host cooperation

THE NEED FOR A 360-DEGREE VIEW OF DATA

Financial institutions need a single source of truth across asset classes, and they also need to support ongoing change and innovation. Because regulators now look at financial services organizations as a single entity, these organizations can't meet regulatory requirements the old way, looking at single-asset classes.

Just as regulators are taking a 360-degree look at financial services organizations for compliance and security, those companies need a 360-degree view of data across the entire organization.

Traditional technologies and approaches make it hard to get that complete view. Companies must combine data across multiple systems and must be ready to perform up-front data modeling and ongoing change control. People expect the process to be slow, expensive, and risky.

DELIVERING INCREASED TRANSPARENCY

Transparency has become increasingly important to regulatory agencies. In order to survive, financial services companies must meet the demands for transparency. Data and processes must be auditable and verifiable.

To meet these demands, financial services organizations must adopt database platforms that improve agility and provide transparency while reducing the costs associated with meeting ever-changing regulations.

“ Reactive compliance needs to become proactive; however, legacy data center infrastructure is the biggest inhibitor to innovation. Financial institutions must modernize their data centers in order to effectively derive insights from both structured and unstructured data, and stay ahead of compliance and regulatory requirements.

– Mike Blalock, General Manager Financial Services Industry, Intel

THE FUNCTIONAL SOLUTION SILOS & TREADMILL

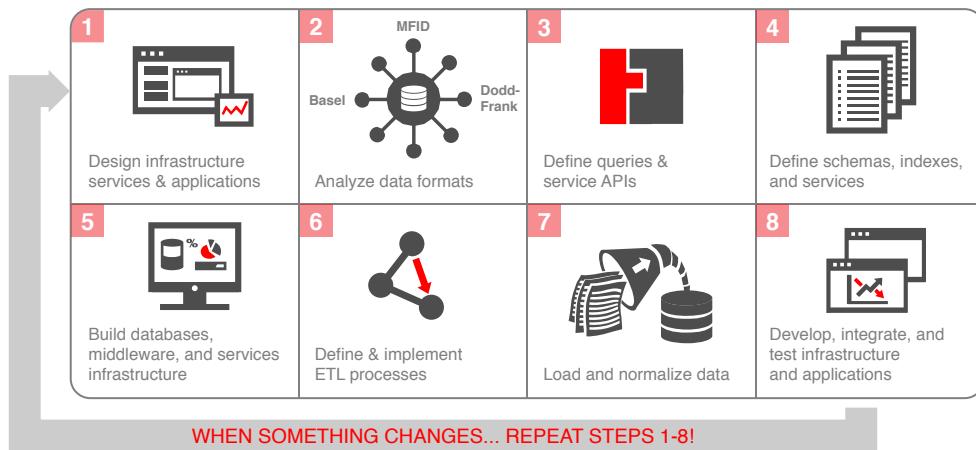


Figure 3

Organizations need to build a secure, transactional, and operational ACID-compliant platform for reporting, applications, and analytics. It's time to leverage insights across the entire organization's data to reduce risk, increase transparency, and demonstrate compliance in an increasingly complex regulatory environment.

ADAPT TO AN EVOLVING REGULATORY LANDSCAPE

Following the financial crisis of 2008, laws such as the European Union's European Market Infrastructure Regulation and the United States' Dodd-Frank Wall Street Reform and Consumer Protection Act were introduced in an effort to reduce the "systemic risk" a bank poses to the financial system. Global financial institutions must now measure how exposed they are to possible losses and then mitigate risks by increasing capital reserves and/or restructuring their portfolio. Data preparedness will determine how smoothly this process goes.

In order to measure a bank's exposure, common calculations must be applied to all trading activities – but reconciling information across the data silos within an organization is a daunting challenge.

DATA AGGREGATION AND REPORTING PRINCIPLES

The Basel Committee on Banking Supervision (BCBS) has developed regulatory standards – essentially, guidelines and principles – to help financial institutions comply with the new laws. In particular, regulation number 239, entitled Principles for Effective Risk Data Aggregation and Risk Reporting, has four core principles:

- **Risk-data aggregation capabilities.** Banks should develop and maintain strong risk-data aggregation capabilities to ensure that risk-management reports reflect the risks in a reliable way.
- **Overarching governance and infrastructure.** A bank should have in place a strong governance framework, risk-data architecture, and IT infrastructure.
- **Risk-reporting practices.** The content of risk reports should be accurate, clear, and complete. Reports should be presented to the appropriate decision makers in time to allow for an appropriate response.
- **Supervisory review, tools, and cooperation.** Supervisors should determine whether the principles achieve the desired objectives.

IMPROVED COMPLIANCE AND LOWER COST

While data integration is critical to delivering a 360-degree view that meets the demands of today, the challenges around data integration with relational databases and ETL make data integration difficult and expensive. Projects can take so long to complete and be so costly that there is essentially an “integration tax” for relying on relational solutions to data integration.

Financial organizations everywhere are seeking ways to be faster and more comprehensive in their approach to integrating data, and organizations that find solutions that avoid the integration tax will find increased productivity, profitability, and adaptability – along with greatly improved compliance.

THE TREADMILL OF ETL

In the past – and even today – many organizations used point solutions that created data silos and applications to access the data within a single silo. As the need grew to integrate data across silos, companies started building complex ETL processes to extract data from the relational databases or other data stores across the company, analyze that data, and load it into yet another set of databases.

This ETL process is not just time-consuming but costly. Change control processes mean that a change in any one of the data sources requires reengineering the ETL process. It's a treadmill that hampers adaptability and productivity. As regulations change across the world, the ETL treadmill puts companies at risk of being unable to meet rapidly evolving compliance requirements. The MarkLogic Enterprise NoSQL platform, optimized on industry-leading Intel® Xeon® processors, provides a single, unified platform with a database, built-in search capabilities, and application services. This reduces the up-front work needed, speeds data integration, and enables sustainable evolution of your data – without complex ETL or data normalization processes – and gets you off the ETL treadmill.

MARKLOGIC AND INTEL® ENABLE FLEXIBILITY, SPEED, AND COMPLIANCE

With all the challenges facing financial services organizations today, it's vital that you establish a sound application framework that will enable efficient, timely response to regulators. The volume and variety of data to be processed and analyzed, the complexity of the analytical processes, and the urgency of the reporting requirements all create technology challenges in the data center.



Figure 4

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.

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

SCALABLE PROCESSORS DESIGNED FOR DATA CENTERS

An example of this is the Xeon E5-2600 v4 product family, which is designed for workloads of a data center, including numerous enhancements that provide performance increases up to 2.21 times that of the Xeon E5-2600 v2 product family, world-class energy efficiency, and enhanced security.

INTELLIGENT AND EFFICIENT STORAGE

The Intel® SSD DC 3700 series solid state drives (SSD) deliver high performance and low latency. Nonvolatile 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.

Intel SSD DC P3700 in combination with the new Intel Xeon processor E5-2699 v4 reduced the data ingestion time, enabled faster reindexing, and delivered significantly higher number of transactions per second, compared to previous generations.

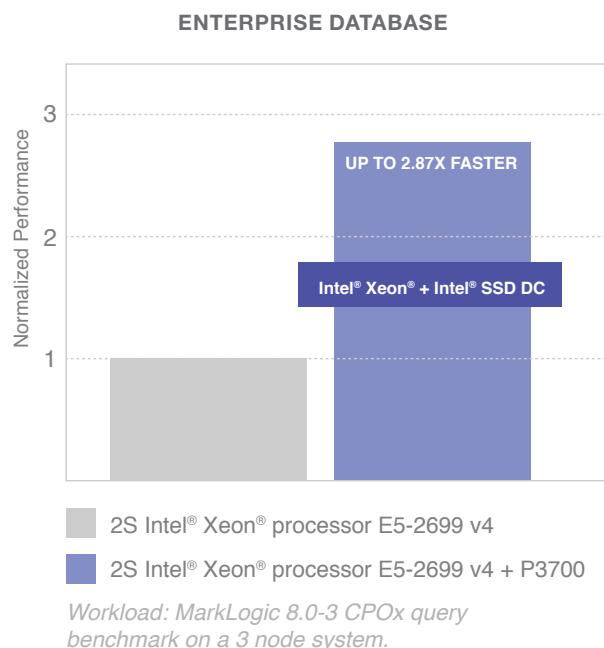


Figure 5 - MarkLogic NoSQL increased performance by up to 2.87X with the Intel® S-SD IDC P3700.

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 association.

The challenges financial organizations face today are not addressed by either hardware or software in isolation but by the harmonious cooperation among all technology components: processors, data storage (including solid-state drives), networking components, database management system software, and business logic/application code. To quote the January 2015 BIS BCBS239 Adoption progress report, “A bank should design, build, and maintain data architecture and IT infrastructure which fully supports its risk data aggregation capabilities and risk reporting practices not only in normal times but also during times of stress or crisis, while still meeting the other principles.”

In order to meet the challenges the financial industry faces, Intel invests considerable resources in working with software providers to assist them in optimizing their applications for Intel technology. In doing so, Intel helps ensure interoperability and solid integration among all the technology components from hardware to software.

SOLUTION ARCHITECTURE FOR REGULATORY REPORTING

Figure 6 depicts the solution architecture for a complete regulatory reporting solution, based on the MarkLogic Enterprise NoSQL database platform using Intel Xeon processor-based servers. Data from upstream trading systems is collected and ingested, via an event stream processing component, into a trade repository, which contains not only trade data but also data about trading counterparties and reference data. Legal documentation in XML format may also be stored to satisfy regulatory demands to correlate contractual agreements with trades executed under such contracts.

This creates a data storage requirement, which for some large banks will exceed a petabyte of data.

Data from the trade repository is then processed and analyzed for downstream reporting to the various regulatory bodies. Two key features of the downstream reporting are:

- **Regulatory store:** Historically, banks have created reports for the regulators but have not stored them for future reference or back-processing. A key feature of this architecture is that all reports generated are stored in a report repository so they can be back-referenced.
- **XBRL (Extensible Business Reporting Language):** This has now been adopted by many of the regulatory bodies for reporting purposes. A key feature of the solution architecture is the ability to generate the reports in the standardized form of XBRL for flexibility and extensibility.

An agile regulatory submission platform facilitates the production of reports from the regulatory store while the submission gateway facilitates the transmission of those reports, in XBRL and other formats, to the relevant regulators.

Analytics capabilities may be provided by SAS or by R from Revolution Analytics, while Qlik and Tableau may be used for more basic regulatory reporting requirements.

Several features of the MarkLogic Enterprise NoSQL database platform running on Intel Xeon-based servers are also worth emphasizing.

MarkLogic Semantics provides a new approach to modeling data that focuses on relationships and context. This is particularly relevant for the correlation of data, such as legal documentation and trade data. It also extends MarkLogic's built-in search capability, providing the ability to expand searches to include related terms or to show users the connection among related entities.

The MarkLogic Bitemporal feature allows the querying of data across system and valid time axes. This is of particular importance for regulatory requirements. It avoids the increasingly harsh consequences from not adhering to government and industry regulations and for audits. Bitemporality preserves the history of all data, including the changes made to it, so that clear audits can be conducted without having to worry about lost data, data integrity, or cumbersome ETL processes with archived data.

MarkLogic Enterprise NoSQL uses Intel® Advanced Encryption Standard New Instructions (AES-NI) to boost cryptography performance for data in all states. From processing to storage to transmission, AES-NI improves data security without sacrificing performance.

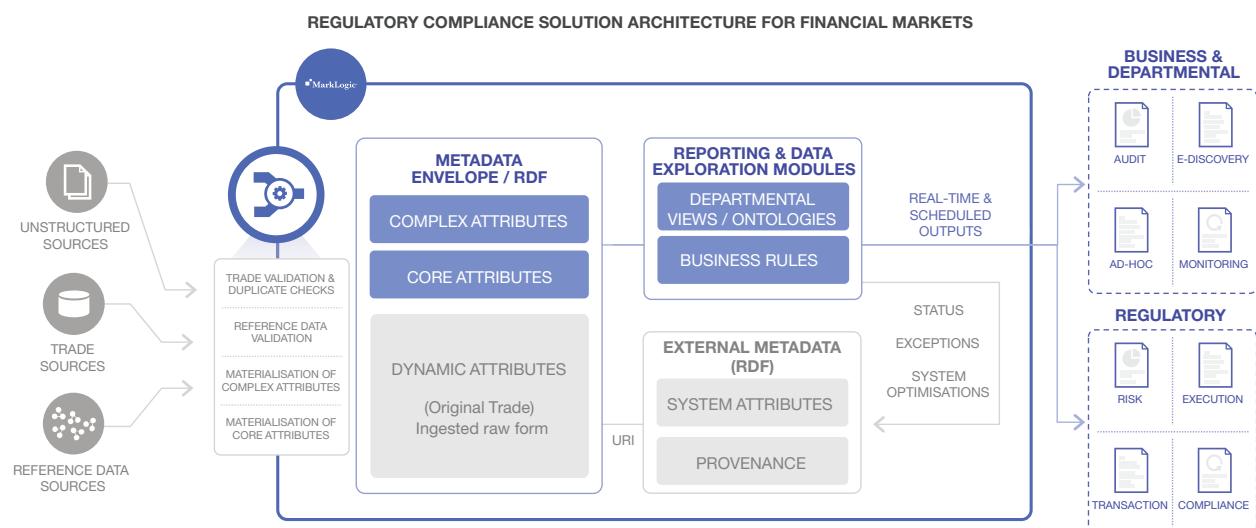


Figure 6



Figure 7 - By instituting MarkLogic, complex ETL can be eliminated from regulatory and compliance process. Business intelligence, strategy, and line of business become streamlined as well.

CASE STUDY: ADAPTING TO REGULATORY COMPLEXITY

The derivatives business of a leading global financial services firm faced problems adapting to the complexity of new regulations and real-time supervisory requirements. The existing relational database management system (RDBMS) platform for its \$70 trillion derivatives business caused poor performance and high cost.

Each derivative processing function was stored in a separate database, resulting in more than 20 relational databases. Derivative trade data came in heterogeneous XML documents, which were difficult to decompose into relational tables. As a result, the organization relied on complex ETL processes, and analysts had to look through more than a dozen different applications to get a comprehensive view of data.

MarkLogic's native XML support streamlined workflow processes and simplified much of the existing architecture, eliminating ETL delays and enabling hundreds of users to perform both structured and ad hoc queries over trade data. With MarkLogic's flexible data model, it is much easier for the back office to accommodate changes to the structure of trade documents without downtime or schema changes. The operational data store provides full ACID compliance and government-grade security to ensure data reliability, integrity, and consistency. In addition, it integrates with a number of other applications to ensure a streamlined trade settlement process.

The result is a successful and growing deployment that powers more than \$70 trillion in trades, persists more than a billion documents, and supports more than 10,000 reads and updates per second. This solution is only one example of how MarkLogic Enterprise NoSQL and Intel Xeon-powered servers provide scalability, adaptability, and agility to handle rapidly evolving data.

“ Only MarkLogic provides a unique multi-model operational and transactional database that can help you build a dynamic GRC strategy. We empower organizations to help expedite regulatory compliance, provide a 360-degree view of enterprise risk, and enable better data governance, all at less cost.

– Ken Krupa, Enterprise CTO, MarkLogic

SUMMARY: MARKLOGIC AND INTEL PROVIDE ADAPTABILITY AND PROACTIVE COMPLIANCE

To be ready for whatever the future holds, organizations need a unified, actionable, 360-degree view of data. Unfortunately, data is spread across disconnected databases. Data integration drags down the speed of business. The underlying problem is that the traditional approach to data integration using relational databases with ETL sacrifices agility, timeliness, and cost.

MarkLogic is the world's best database for integrating data from silos. Organizations around the world rely on MarkLogic—an operational and transactional Enterprise NoSQL database platform—to integrate their most critical data and build innovative applications on a 360-degree view. MarkLogic makes it easy to get data in and easy to get data out, and it is 100 percent trusted to run critical business operations.

ABOUT MARKLOGIC

Today, more data is collected than ever before, and organizations have great ideas for bigger, smarter applications they can roll out to millions of users. Unfortunately, in today's world of big data, the relational databases that have been relied on for the past few decades are too limiting and inflexible. Organizations are facing a growing inability to handle their current structured data, let alone the new data sources that go unaccounted for.

Since our inception in 2001, MarkLogic Corporation has focused on offering organizations a new-generation database platform designed to integrate, store, manage, and search more data than ever before. Global enterprises and governments rely on MarkLogic to power intelligent, high-performance applications aimed at analyzing data for better insights and also running crucial day-to-day business operations.

With MarkLogic, organizations achieve faster time-to-value than was possible with legacy relational databases, while minimizing risk and improving data quality.

For more information about MarkLogic, visit: www.marklogic.com.

ABOUT INTEL

Intel (NASDAQ: INTC) expands the boundaries of technology to make the most amazing experiences possible. Intel powers the incredible technology at the heart of transforming any business or industry through powering supercomputers and the cloud, and providing the tools and resources that will help organizations evolve.

Additional information about Intel is available at www.intel.com.

For more information about Intel in Financial Services visit: www.intel.com/financialservices.

© 2017 MARKLOGIC CORPORATION. ALL RIGHTS RESERVED. This technology is protected by U.S. Patent No. 7,127,469B2, U.S. Patent No. 7,171,404B2, U.S. Patent No. 7,756,858 B2, and U.S. Patent No 7,962,474 B2. MarkLogic is a trademark or registered trademark of MarkLogic Corporation in the United States and/or other countries. All other trademarks mentioned are the property of their respective owners.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Intel, the Intel logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© 2017 Intel Corporation.

MARKLOGIC CORPORATION

999 Skyway Road, Suite 200 San Carlos, CA 94070

+1 650 655 2300 | +1 877 992 8885 | www.marklogic.com | sales@marklogic.com



999 Skyway Road, Suite 200 San Carlos, CA 94070
+1 650 655 2300 | +1 877 992 8885
www.marklogic.com | sales@marklogic.com