

REGULATORY REPORTING FOR FINANCIAL SERVICES

Gordon Hughes, Global Sales Director, Intel Corporation
Sinan Baskan, Solutions Director, Financial Services, MarkLogic Corporation



Many regulators and regulations

Financial Services industry is under constant pressure from financial regulation bodies to ensure the integrity of the system

- Prudential Regulation Authority
- Financial Conduct Authority
- FINREP
- COREP
- Basel Accord
- Sarbanes-Oxley Act
- Reporting required daily, monthly, quarterly, annually
- Volume and variety of data is vast
 - Trades executed
 - Reference and counterparty data

The Imperative – Regulation and Restructuring

“ *Data architecture and IT infrastructure – 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. ”*

“ *A bank should establish integrated data **taxonomies** and architecture across the banking group, which includes information on the characteristics of the data (**metadata**), as well as use of single identifiers and/or unified naming conventions for data including legal entities, counterparties, customers and accounts.*

Roles and responsibilities should be established as they relate to the ownership and quality of risk data and information for both the business and IT functions. ”

Basel Committee on Banking Supervision
Principles for effective risk data aggregation and risk reporting, 2013

And the Answer is Out of the Box

Banks do not necessarily need to have one data model; rather, there should be robust automated reconciliation procedures where multiple models are in use

*Principles of Effective Risk Data Aggregation and Risk Reporting
Basel Committee on Banking Supervision, BIS, January 2013*

*... “all kinds of data, structured and unstructured, internal and external, are gathered. **Data does not need to follow strict rules when it enters the bank** (as it does in a data warehouse).*

*Rather the user of data defines the rules when extracting data from the lake. Combined with Google-like **search technology, the data lake enables a step-change for banks to leverage their data** for all purposes such as marketing, risk and finance.”*

McKinsey&Company


Who is looking for What? What's Valuable to Them?

CEO - Growth



- Comply with Regulations
- Increase margins by lowering cost of sales and support services
- Improve customer loyalty and brand
- Improve upsell/cross-sell for revenue growth
- Deliver Shareholder Value (ROE)

COO - Execution



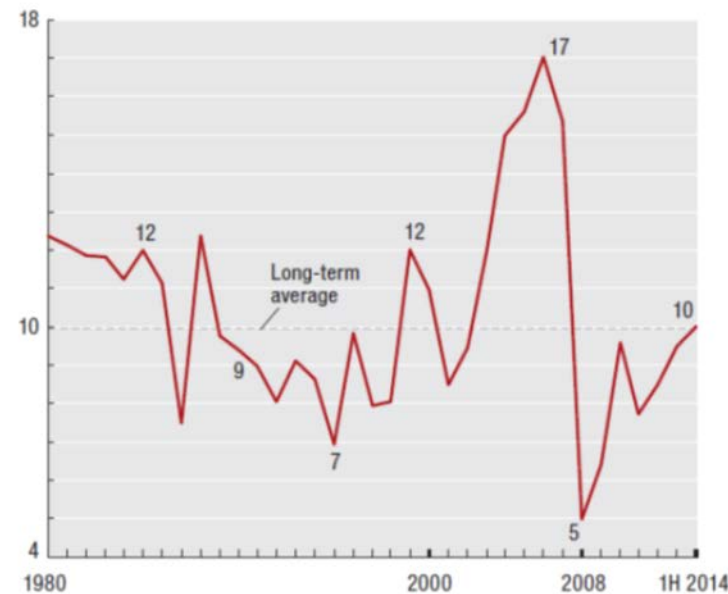
CIO - Enablement



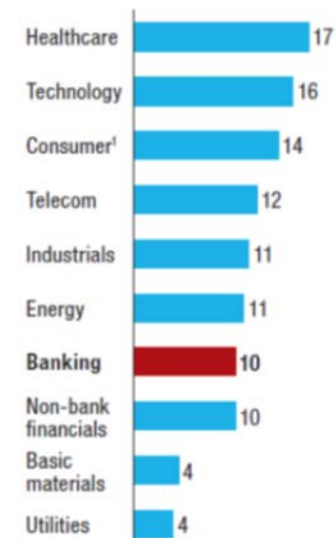
CFO - Risk



Global banking ROE, 1980-1H 2014
Percent



ROE, 2013
Percent



¹ Cyclicals and non-cyclicals
Source: Bloomberg; Compustat; Datastream; OECD; Reuters; McKinsey Panorama—Global Banking Pools

Clarion Call

Regulations

- Regulations and reform drive re-structuring and strategic transformation

Invest in Transparency and visibility

Run the Business

- Consolidation and global expansion compel operational excellence

Manage complexity, dis-intermediate silos

Disruption vs Innovation

- New Competitors arriving and next generation technologies at lower entry points

Manage out of “legacy” limitations; innovate w/ product & service

Customers

- Know Your Customer + Understand Your customer

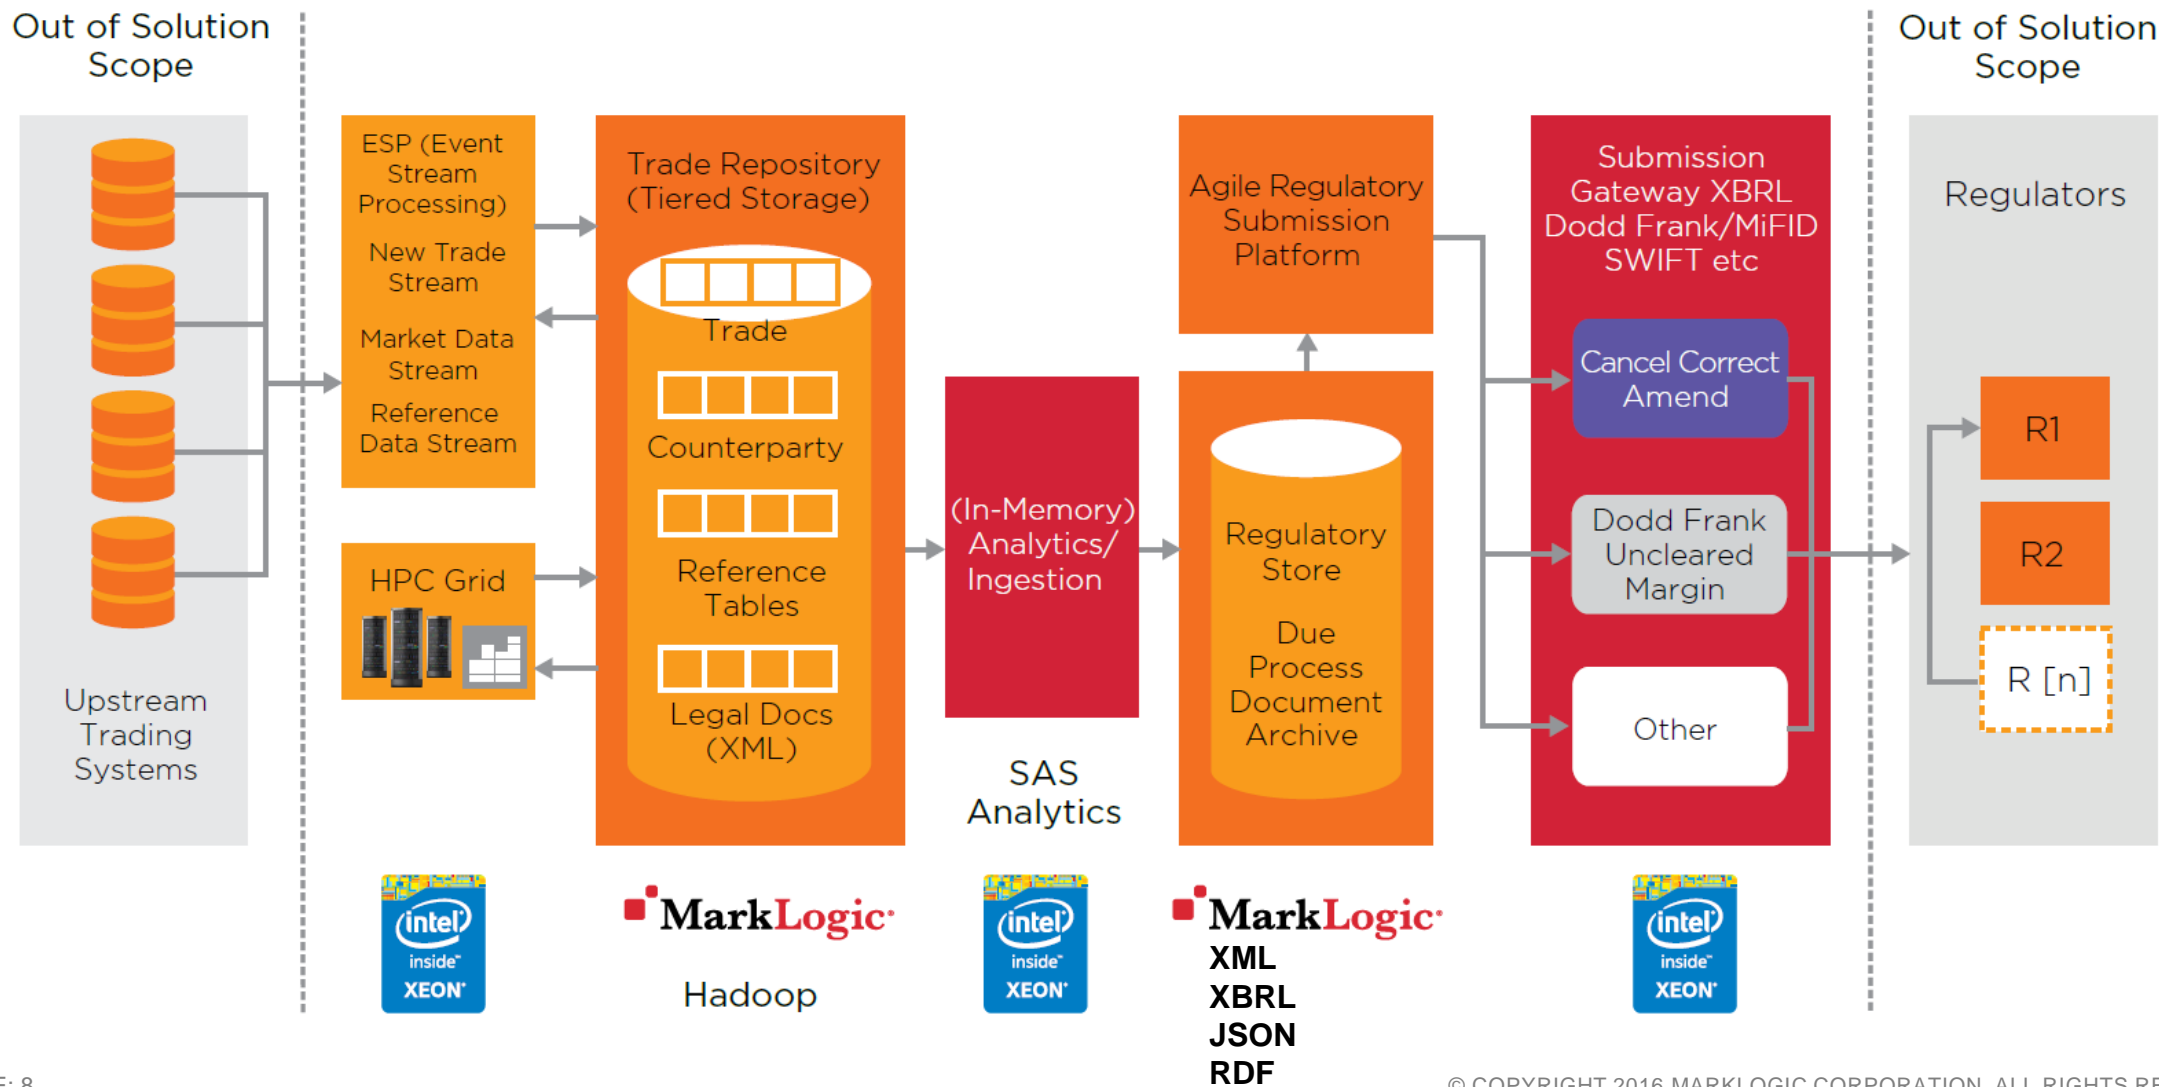
Customer 360 view; cross-product, self-service

Solution Architecture for Regulatory Reporting

Complete MarkLogic and Intel Solution

- Data from upstream trading systems is collected and ingested into a trade repository
 - Legal documentation in XML format may also be (can exceed one petabyte)
- Data from the trade repository is then processed and analyzed for downstream reporting to the various regulatory bodies
 - Regulatory Store
 - XBRL (Extensible Business Reporting Language)

Regulatory Reporting Process & Architecture



Intel® Xeon® Processor E5-2600 v4 Product Family

MarkLogic*

NoSQL for the Enterprise

“The Intel® Xeon® processor E5-2699 v4 boosted query performance compared to previous generations processors, providing great benefits to query-intensive workloads, while the Intel® Solid State Drive DC P3700 improved ingestion and re-indexing performance to amazing levels compared to traditional HDDs, enabling new transaction speed levels.”¹

Caio Milani – Director of Product Management, MarkLogic

- MarkLogic is the only Enterprise NoSQL database, ideally suited for making heterogeneous data integration simpler and faster and for doing dynamic content delivery at massive scale.
- The combination of Intel® SSD DC P3700 and the new Intel Xeon processor E5-2699 v4 reduced the data ingestion time, enabled faster re-indexing and delivered significantly higher number of transactions per second, compared to previous generations (see the chart).
- The Intel® Xeon® processor E5-2697 v4 delivers about 13%² improvement in query performance compared to E5-2697 v3.


Service more customers concurrently – deliver a better user experience!



www.marklogic.com

Enterprise Database

MarkLogic NoSQL increased performance by up to 2.21X with 2S Intel® Xeon® processor E5-2699 v4



Workload: MarkLogic 8.0-3 CPOx query benchmark on a 3 node system.

¹ - Testing conducted on MarkLogic software comparing 2S Intel® Xeon® Processor E5-2699 v4 to 2S Intel® Xeon® Processor E5-2697 v2 and 2S Intel® Xeon® Processor E5-2697 v3. Testing by MarkLogic. For complete testing configuration details, [SEE SLIDE 120](#). ² - MarkLogic estimated increase for 2S Intel® Xeon® Processor E5-2697 v4 to 2S Intel® Xeon® Processor E5-2697 v3. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit <http://www.intel.com/performance>.

MarkLogic Optimized by Intel Technology

MarkLogic is Designed for Enterprise-Grade Data Integration




SCALE-OUT COMMODITY HARDWARE

Scalable and elastic without over provisioning or over-spending



BITEMPORAL

Time-based history and audit trail of data



TIERED STORAGE

Minimize the cost of storage as database expands



GEOSPATIAL

Integrate search with location and geospatial coordinates



SEMANTICS

Enhance search and define data relationships



REAL-TIME ALERTING

Unlimited alerts on new data at time of ingestion



ACID TRANSACTIONS

Scalable performance and multi-version concurrency control



GOVERNMENT GRADE SECURITY

Certified security and role-based access control at document level

The Massive Cost of Integrating Data From Silos

80 % OF TIME
WASTED


By data scientists just
wrangling data

60 % OF THE
COST

Of data warehouse projects
is on ETL

\$36 BILLION IN
SPENDING


In 2015 on creating relational
data silos



THE IT CHALLENGE

Relational Databases with ETL Sacrifice Agility, Timeliness, and Cost

- All future data needs must be predictable
- New SQL queries require database re-indexing
- Siloed database changes require ETL re-writes



THE DESIRED SOLUTION

**A Database That
Integrates Data Better,
Faster, with Less Cost**

The MarkLogic Alternative

An *Operational and Transactional* Enterprise NoSQL Database



EASY TO GET DATA IN

Flexible Data Model

- Data ingested as is (no ETL)
- Structured and unstructured data
- Data and metadata together
- Adapts to changing data and changing data structures



EASY TO GET DATA OUT

Ask Anything Universal Index

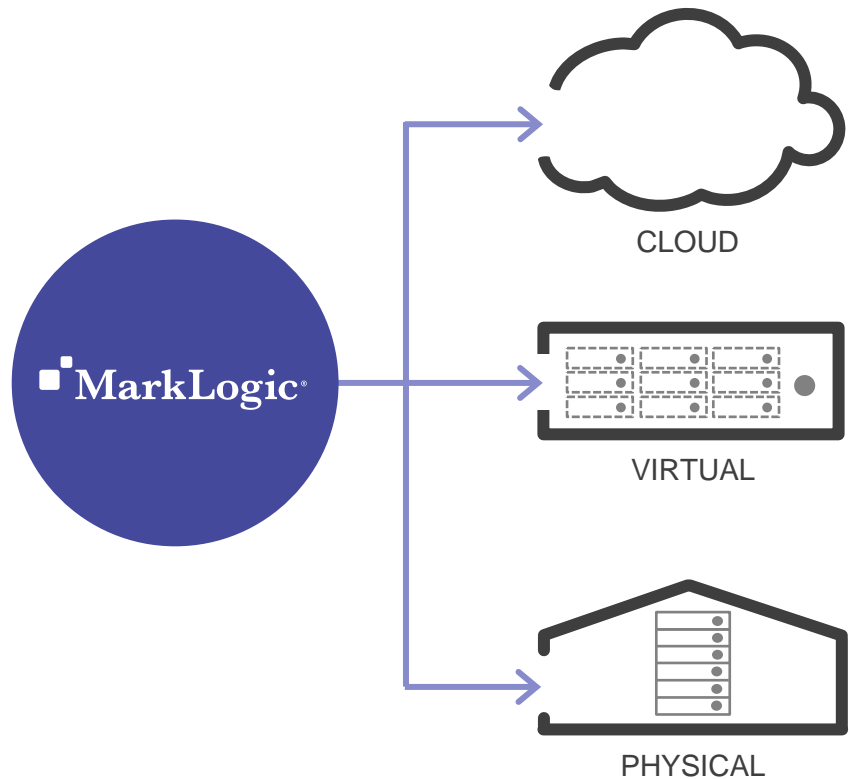
- Index once and query endlessly
- Real-time and lightning fast
- Query across JSON, XML, text, geospatial, and semantic triples in one database



100% TRUSTED

Enterprise Ready

- Reliable data and transactions (100% ACID compliant)
- Out-of-the-box automatic failover, replication, and backup/recovery
- Enterprise-grade security and Common Criteria certified



Flexible Deployment Run It Wherever You Want

- Linux or Windows
- On-premises (virtual or physical)
- Private cloud, hybrid cloud, public cloud
- AWS or Azure
- Any storage architecture, including Hadoop

Why MarkLogic?

VALIDATED

Fast Time to Results



Ask Anything Universal Index



Trusted Data and Transactions



Enterprise-Grade Security



Scale-Out Commodity Hardware



Lightning Fast and Real-Time



Q&A