

Simplifying Data Governance and Accelerating Real-time Big Data Analysis in Financial Services with MarkLogic Server and Intel

Reduce risk and speed time to value using an integrated NoSQL + Big Data solution built on MarkLogic and Intel

Executive Summary

The era of Big Data and enterprise government regulation are driving considerable changes in how Financial Services companies manage and use the varied types of data they acquire and store. The legacy Relational Database Management System (RDBMS), Enterprise Data Warehouse (EDW), and Storage Area Network (SAN) infrastructure used by companies today to create siloed data environments is too rigid to accommodate the demands for massive storage and analyses on a larger and wider variety of data. Forcing this legacy architecture into today's enterprise requirements is costly and risky.

MarkLogic has integrated their new-generation Enterprise NoSQL platform with Apache Hadoop* optimized for Intel® Architecture to deliver a powerful platform ideal for Financial Services. This solution provides all the features of MarkLogic and Apache Hadoop with the data governance and security IT departments need. Running on Intel® technology and the enhancements Intel has brought to Apache Hadoop, this integration gives companies a true Big Data solution with enterprise-class security for storage, real-time queries, and analysis of all their data.

This paper summarizes the issues companies face today with legacy RDBMS + SAN data environments and why the combination of MarkLogic, Apache Hadoop, and Intel provides a solution for enterprise Big Data.

“MarkLogic has integrated their new-generation Enterprise NoSQL platform with Apache Hadoop* optimized for Intel® Architecture to deliver a powerful platform ideal for Financial Services.”

Table of Contents

Executive Summary 1

MarkLogic, Hadoop, and Intel in
Financial Services 2

The Criticality of Today's Data
Governance 3

 The Rigidity of Traditional
 Enterprise Data Environments. . 3

 Unstructured Data Drives
 Change 3

 MarkLogic Enterprise NoSQL... 3

The Era of Big Data and Apache
Hadoop* 4

 MarkLogic with Apache
 Hadoop 5

 Apache Hadoop 5

MarkLogic and Intel 5

MarkLogic, Hadoop and Intel in
the Field 6

BBD—Before Big Data 6

Summary 7

MarkLogic on Intel® Technology .. 7

MarkLogic, Hadoop, and Intel in Financial Services

The world's largest financial institutions have adopted MarkLogic's enterprise NoSQL database to aggregate disparate data types from the myriad of systems that span their enterprises. MarkLogic quickly loads data "as-is," eliminating months of development work associated with relational EDWs and ETL. The agility enabled by MarkLogic is helping banks deal with cost pressures, customer needs, heightened regulatory scrutiny, and demands for greater transparency.

Below are some examples of the pain points that MarkLogic is helping financial institutions address (organized by use cases):

Operational Trade Data Store Modernization

- Quickly operationalize new instruments in response to market conditions
- Comply with new regulations requiring trade reconstruction
- Minimize trade exceptions and cost per trade
- Reference Data Management
- Quickly address changing requirements of data consumers
- Quickly respond to changes on the data vendor side
- Support bi-temporal views and provide data provenance required by regulators

Customer Insight

- Integrated data from multiple silos and Lines of Business without extensive modeling and ETL efforts
- Obtain a 360-degree customer view to support marketing, regulatory, and risk management needs
- Tap into valuable information locked in non-relational sources, both inside (e.g., onboarding docs) and outside the company (e.g., social media)

Fraud, AML, and Financial Crime Prevention

- Rapidly integrate new information from drastically diverse sources (transactions, customers, employees, and devices) as it becomes available
- Transition from a transaction-based detection to a customer-centric detection to provide a more holistic view and minimize false positives
- Respond to suspicious activities before fraud/crime is committed
- Incorporated data from non-relational sources

Decision Support and Information Data Platform

- Support multiple delivery mechanisms, including subscription-based and alerting
- Handle multiple types of data and analytics, including geospatial, text, and sentiment
- Quickly integrate new data sources into a single, coherent information delivery platform

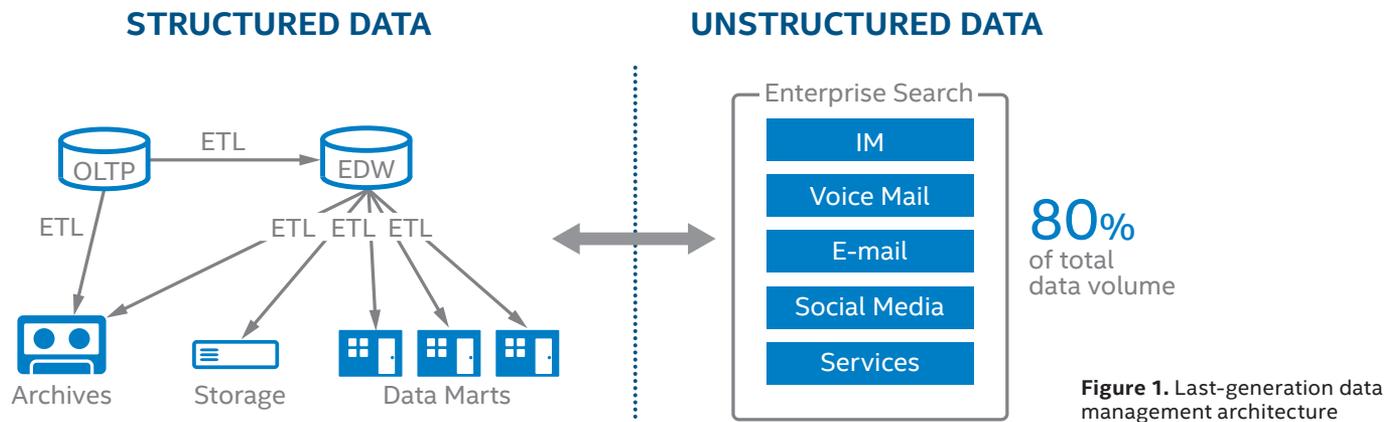


Figure 1. Last-generation data management architecture

The Criticality of Today's Data Governance

What drives business success is actionable knowledge from data. The extent and value of that data—from daily transactions to emails, text messages, contracts, and others—is more critical today than ever. And, with new rules affecting regulated industries, the policies and processes to capture, manage, and protect that data has come under scrutiny from government auditors. These impacts are driving enterprises to take a new look at how they deal with data.

The Rigidity of Traditional Enterprise Data Environments

Companies have long used online transaction processing (OLTP) systems based on Relational Database Management Systems (RDBMS) plus Storage Area Network (SAN) to gather the essence of their daily activities, which analytical systems periodically process. Company growth and queries beyond what the schema of the RDBMS was designed to provide, however, eventually result in a system that no longer serves the wide-ranging needs of the organization. These effects potentially cause lost revenue, missed opportunities, and more.

To adapt to information demands within the enterprise, IT often spins off enter-

prise data warehouses (EDW) from the company's RDBMS to create dedicated report and analytical systems that serve a specific application (Figure 1). These siloed data environments require more investment, more Extract-Transfer-Load (ETL) operations, and duplicated data, creating greater burden for IT, rising costs, and increased risk.

Eventually, more creative or complex analyses are required that the EDW cannot provide. Thus, individual departments within the organization create smaller data marts extracted from the EDW. These data marts provide the core content for real-time analysis using Excel* and other business intelligence tools. The results are even more copies of data pools and information systems potentially beyond IT's visibility and manageability.

At some point, scaling up the database, EDW, data marts, and storage for these proprietary systems becomes economically unsustainable. In order to maintain predictability in cost and performance of the most important data, companies archive the less important (usually older) data. But, finding the right slice of data across a large RDBMS schema is challenging. Again, it introduces brittle and costly ETL, and the archives are unavailable for deep analytics that might be needed. Additionally, the data marts unknown to IT might lag, working

with older and possibly inaccurate information.

Unstructured Data Drives Change

This evolution has been repeated across enterprises in order to stay competitive. It challenges any company's data governance capabilities. And, it does not even include schema-less or "unstructured" data.

All of the text contracts, voicemail, email, web clicks, and instant messages that provide the "long tail" of information involving a transaction are simply not captured. Yet, today's legislation requires some institutions, such as banks, to demonstrate a holistic view of activities, such as the trade lifecycle for over-the-counter derivatives. Thus, unstructured data must be part of the enterprise-wide data management.

NoSQL data platforms and the rise of Big Data solutions help address these critical governance issues, in addition to providing new capabilities to gain insight from more varieties of data.

MarkLogic Enterprise NoSQL

For more than a decade, MarkLogic has delivered a powerful, agile, and trusted enterprise-grade, schema-agnostic Not-Only SQL (NoSQL) database platform. Using MarkLogic Server, an organization's entire structured and

unstructured data repository can be stored in one indexed location, enabling fast application building on top of it and eliminating the need for siloed systems. MarkLogic can also be used in a data virtualization/Logical Data Warehouse environment to bridge data silos that cannot be merged in one location for whatever reason. This approach allows organizations to more quickly turn all data into valuable and actionable information in real time, while reducing risk, cost, and management overhead.

Among its key features, the MarkLogic platform includes:

- Atomicity/Consistency/Isolation/Durability (ACID) transactions
- Horizontal scaling and elasticity
- Real-time indexing
- Search and semantics
- High availability
- Disaster recovery
- Replication
- Government-grade security

Enterprises around the world and across a wide range of industries, including healthcare, entertainment, financial services, retail, government agencies, and others, have adopted the MarkLogic platform to manage and

analyze all their data. These organizations are benefiting from building value from their data instead of schemas and infrastructure to support and understand it.

More recently, the emergence of Apache Hadoop* has brought yet more capabilities for enterprise data storage and analysis. MarkLogic integrates readily with Hadoop in a number of important ways.

The Era of Big Data and Apache Hadoop*

To take advantage of the value of all their data, organizations are aggressively moving toward storing and maintaining data that might have been previously discarded. Saving this legacy information creates a sandbox for data science, enabling possibilities of new and deeper population-level analyses, as well as more traditional data preparation and aggregation (ETL). But, how does a company operationalize this rich information?

As we have seen, the RDBMS model, while still offering the enterprise capabilities organizations have come to expect, constrains what can be done with the data. Business needs a new paradigm.

“For more than a decade, MarkLogic has delivered a powerful, agile, and trusted enterprise-grade, schema-agnostic Not-Only SQL (NoSQL) database platform.”

Apache Hadoop has emerged as a cost-effective place to store raw, intermediate, and finished data of all types—both structured and unstructured (Figure 2). It can accommodate massive amounts of data in any shape—and do it cheaply.

Hadoop also integrates tools for distributed computation across petabyte-sized volumes that are beyond what RDBMS + SAN implementations can do. Apache Hadoop has become the core of Big Data solutions with its staging, persistence and analytics capabilities:

- **Staging:** Load raw data into Hadoop. Use MapReduce* operations to prepare data for other uses, including filtering, aggregation, mash-up, transformation, etc.

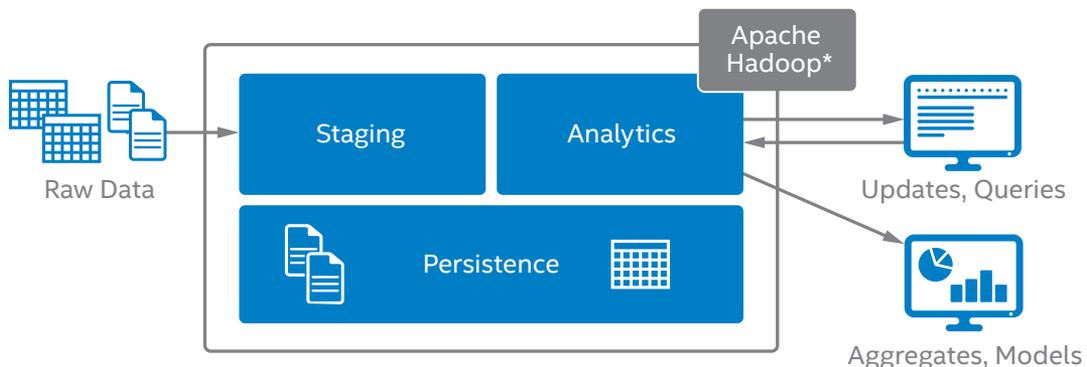
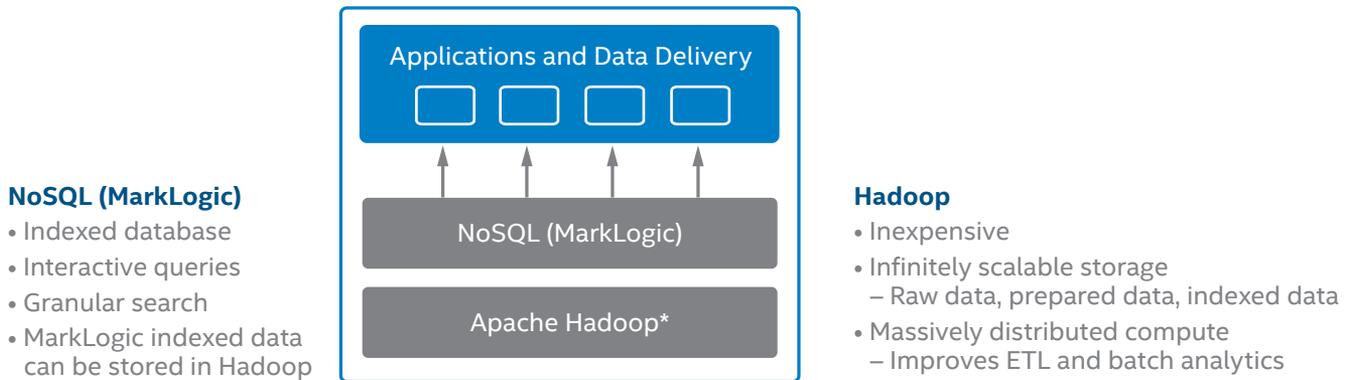


Figure 2. The capabilities of Apache Hadoop*



NoSQL (MarkLogic)

- Indexed database
- Interactive queries
- Granular search
- MarkLogic indexed data can be stored in Hadoop

Hadoop

- Inexpensive
- Infinitely scalable storage
 - Raw data, prepared data, indexed data
- Massively distributed compute
 - Improves ETL and batch analytics

Figure 3. Enterprise NoSQL + Apache Hadoop*: new generation

- **Persistence:** Keep the raw inputs around for later inspirational integration and analytics, without losing the original context. Keep the intermediate prepared data around, also. Manage raw and prepared indexes under the same infrastructure and with the same governance policies.

- **Analytics:** Perform large-scale, population-level analyses on raw or prepared data.

However, while open source Apache Hadoop offers analytics and storage capabilities businesses want today, it was not designed for the real-time applications nor the data governance requirements enterprises need.

MarkLogic with Apache Hadoop

To integrate with Hadoop, MarkLogic enhanced the MarkLogic platform to utilize Hadoop Distributed File System (HDFS) storage (Figure 3). Enterprises can run the MarkLogic database on top of HDFS, providing role-based security, full-text search, ACID transactions, and the flexibility of a granular document data model for real-time applications—all within the existing Hadoop infrastructure.

With MarkLogic's data files stored in HDFS, analysts can also run MapReduce jobs on those files directly. This opens up MarkLogic's formerly proprietary data format to other workloads and makes the file format a viable long-term archive option.

With MarkLogic's indexes stored in HDFS, companies can quickly gather ad hoc subsets of indexed data and attach them to a MarkLogic database to have that data immediately available for interactive updates and queries. This simplifies operations and data governance, maintains the security and metadata when it was first indexed, and allows use of those initial indexes (and security and metadata) throughout the life of the data.

Apache Hadoop

Proven in production at some of the most demanding enterprise deployments in the world, Apache Hadoop is supported by experts at Intel with deep optimization experience in the Apache Hadoop software stack as well as knowledge of the underlying processor, storage, and networking components.

Intel's enhancements to Hadoop are designed to enable the widest range of use cases on Hadoop by delivering the performance and security that enterprises demand. Intel delivers platform innovation in open source, and it is committed to supporting the Apache developer community with code and collaboration.

The combination of MarkLogic, Intel technologies, and Apache Hadoop optimized for Intel Architecture delivers the best of MarkLogic's platform performance, security, and manageability. While a NoSQL + Hadoop solution helps bridge traditional RDBMS with the wider-ranging data analytics and storage capabilities of Hadoop, the combination of MarkLogic and Apache Hadoop enhanced on Intel Architecture makes this an enterprise-class Big Data solution.

MarkLogic and Intel

The combination of MarkLogic and Apache Hadoop enables enterprises to implement both granular, real-time analyses plus deep batch analytics on massive data sets with enterprise-grade data governance—all on top of a single repository (Figure 4). With MarkLogic and Apache Hadoop, rather than building a new dedicated silo of storage, database, warehouse, middle-ware, and thick client, companies can focus on the value of the data instead of the infrastructure and still be assured there is no compromise on performance, availability, or security.

MarkLogic provides the secure, reliable, and high-performance real-time indexing, search, and analysis platform the company’s customers have come to trust. Enhancements based on Intel® technologies offer hardware-enhanced, secure Apache Hadoop operations with significant performance improvements over pure open source Hadoop.

The combination of MarkLogic and Intel dramatically shrinks the applica-tion stack, making building applications much less expensive and less risky. Organizations can more freely innovate and cut their losses early if an idea doesn’t work.

“The combination of MarkLogic and Apache Hadoop delivers the best of MarkLogic’s new-generation platform and Intel’s hardware-enhanced Apache Hadoop performance, security, and manageability.”

BBD—Before Big Data

Only in the last few years have unstructured data, Big Data, and Apache Hadoop* become important parts of enterprise operations. But, before the 2009 release of Apache Hadoop, in 2003 MarkLogic released its Not-Only SQL (NoSQL) database, search engine, and application software platform to enable analytics on both structured and unstructured data (Figure 5). The capability of a data storage, query, and analysis platform beyond Relational Database Management Systems (RDBMS) was born—long before the idea of Big Data. This new-generation MarkLogic platform became the foundation of systems that have given organi-zations the ability to gain insight from and act on more of their data in new ways.

With the emergence and growing adoption of Hadoop across indus-tries and the Big Data storage and processing benefits it offers, MarkLogic integrated Apache Hadoop into the MarkLogic platform (NoSQL + Hadoop). The combination gives Hadoop the real-time search capabilities and enterprise-grade database platform organizations need to operationalize all their data, yet Hadoop still requires additional critical capabilities, like encryption and management tools, IT demands.

MarkLogic has integrated Intel into their system to deliver these needed features—not just NoSQL + Hadoop, but MarkLogic + Intel.

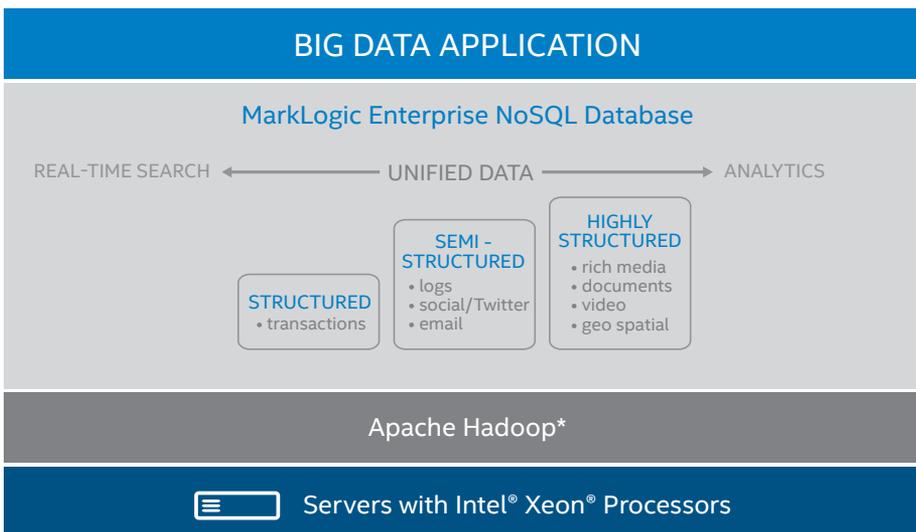


Figure 4. MarkLogic and Apache Hadoop* on Intel® technology



Figure 5. MarkLogic enables 'Big Data' before Big Data

Summary

MarkLogic with Apache Hadoop on Intel Architecture delivers an enterprise-class Big Data solution with the data governance capabilities Financial Service companies need for management and real-time analysis of all their data. The integration of these two solutions creates a platform that helps IT reduce risk and contain—even reduce—cost, while accelerating application development of data analytics to serve the needs of an agile business.

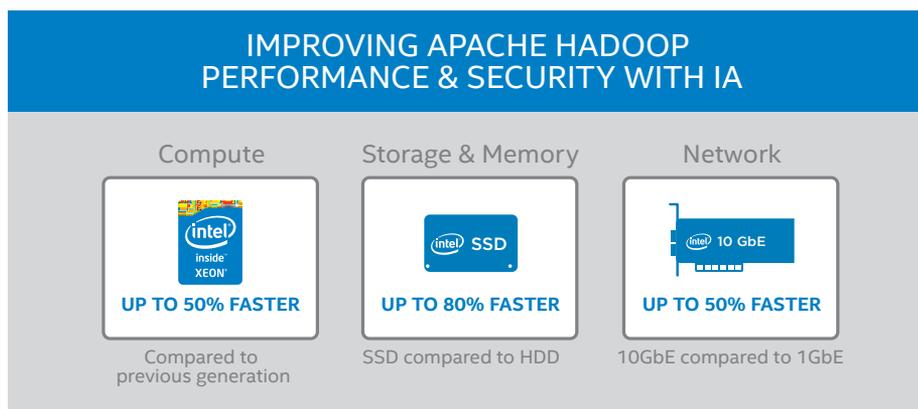
MarkLogic and Intel-enhanced Apache Hadoop allow companies to keep all their data readily available for business intelligence and deeper population-level studies that can provide new insights and reveal new opportunities.

The agility enabled by MarkLogic is helping banks deal with cost pressures, customer needs, heightened regulatory scrutiny, and demands for greater transparency.

For more information on MarkLogic and Intel, visit www.marklogic.com and www.intel.com.

MarkLogic on Intel® Technology

Intel® technology, as a hardware foundation for Apache Hadoop, delivers significant performance improvement for Hadoop processing (Figure 6). Along with Intel® Xeon® processors, Intel® Solid State Drives, and Intel® 10GbE networking, Intel offers a robust contribution to support the new generation of the MarkLogic Enterprise NoSQL platform.



As measured by time to completion of 1TB sort on 10 node cluster

Figure 6. Intel® technology provides accelerated performance for Apache Hadoop*

¹ Hadoop Westmere Test Bed: 4 hours

Hardware Configuration: Arista 7050T; 10 x SuperMicro 1U servers: Intel Processor: 2 x 3.46 GHz Intel[®] Xeon[™] processor 5690; Memory: 48 GB RAM; Storage: 5 x 700 GB 7200 RPM SATA disks; Intel[®] Ethernet 10 Gigabit Server Adapters (10GBASE-T); Intel[®] Ethernet Gigabit Server Adapter (1000BASE-T)

Software Configuration: Operating System: CentOS 6.2; Hadoop: Cloudera's Distribution; Java[®]: Oracle JDK 1.7.0.

Cluster Configuration: 1 Client machine; 1 Head node (Name node, Job Tracker); 9 Workers (data nodes, task trackers).

Network Division Hadoop Romley Test Bed: 7 minutes

Cluster Configuration: 1 Head Node (name node, job tracker); 10 Workers (data nodes, task trackers); 10-Gigabit Switch: Cisco Nexus 5020;

Software Configuration: Intel Distribution for Hadoop 2.1.1; Apache Hadoop 1.0.3; RHEL 6.3; Oracle Java 1.7.0_05.

Head Node Hardware: 1 x Dell r710 1U servers: Intel: 2x3.47GHz Intel[®] Xeon[™] processor X5690; Memory: 48 GB RAM; Storage: 10K SAS HDD; Intel[®] Ethernet 10 Gigabit SFP+; Intel[®] Ethernet 1 Gigabit.

Worker Node Hardware: 10 x Dell r720 2U servers: Intel: 2 x 2.90 GHz Intel[®] Xeon[™] processor E5-2690; Memory: 128 GB RAM; Storage: 520 Series SSDs x 5; Intel[®] Ethernet 10 Gigabit SFP+; Intel[®] Ethernet 1 Gigabit.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

Copyright © 2014 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others. Printed in USA 0614/JO/OCG/PDF ♻️ Please Recycle 330578-002US

