



## Converging Analysis and Operations

WHETHER IT'S understanding customers, managing operational risks, or positioning an enterprise's information assets for competitive advantage, today's business decision makers realize that the 21st century economy is more information-driven than ever before. Yet all too often, line-of-business needs have resulted in application-centric development that keeps valuable enterprise data locked up in business-aligned data silos.

For years, two Enterprise Architecture patterns have in various ways promised to break down enterprise information silos. First there was the Enterprise Data Warehouse (EDW) which brought together disparate line-of-business data in support of decision support analysis. Later came Service-Oriented Architectures (SOA), which promised to orchestrate information exchange across business lines in a standard way. In each case, but for different reasons, there's a focus on moving data from one place to the next. In the case of the data warehouse, the data movement is ETL-related (extract, transform, and load). In the case of SOA, the focus is on passing data between business silos.

For today's information-driven architectures, transforming, moving, and otherwise copying data for different needs has become an inhibitor. Gartner's three "big data" V's of volume, variety, and velocity make an ETL and/or copy-centric approach untenable. As a result, some enterprises have implemented Hadoop in the hopes of addressing these three V's. However, lacking transactional capability, as well as a mature security and governance model, most have fallen short and have had to rely on a Hadoop + RDBMS strategy to fill the gaps. The issue in this case, however, is more data copying and enterprise-level ETL.

### A NEW ARCHITECTURE: THE OPERATIONAL DATA HUB

Forward-thinking enterprises are adopting new approaches that go beyond a discovery-only Data Warehouse or point-to-point SOA data integration. Instead, they align and unify to a **data-centric** architecture.

The MarkLogic® Operational Data Hub solution helps organizations better leverage information assets—by implementing a data-first enterprise architecture for all business data. The Operational Data Hub architecture brings applications to the data, instead of requiring you to spend time and money moving and copying data between applications. It brings architecture agility to your organization, letting you incrementally implement your data strategy—and embrace change as the norm, not only for data analysis but also for the operations that run your enterprise.

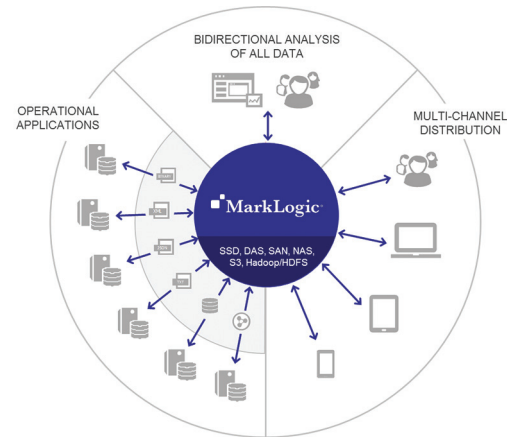
And, because the Operational Data Hub solution is based on the MarkLogic Enterprise NoSQL database platform, you are ensured enterprise-grade reliability, resiliency, consistency, and government-grade security for your information assets.

The flexible, agile Operational Data Hub architecture integrates heterogeneous data and provides your organization with powerful data services to support real-time business applications, delivering:

- Overall reduction in application delivery time—without sacrificing data governance
- A stable platform for evolving a Master Data Management strategy
- Better ability to manage change and innovate as new business needs arise
- A closer alignment of analytical activity to the operational activity it is intended to support

### ADVANTAGES OF ENTERPRISE NOSQL

Organizations who have implemented MarkLogic benefit from a unified data management platform that lets them avoid the complexity and brittleness inherent in integrating multiple systems. MarkLogic is built with a flexible data model to ingest, store, manage, and search today's data, without sacrificing any of the data resiliency and consistency features of last-generation relational databases, enabling organizations to:



- Handle massive volumes of heterogeneous data, metadata, and content in one platform, without the pain of upfront data modeling
- Ensure data integrity and reliable performance with transactional consistency (ACID), high availability, government-grade security, and real-time indexing
- Scale out on any hardware environment, adapt as data changes and new data is added
- Support auditing requirements in highly regulated industries with bitemporal data
- Assert new facts and relationships about data in real time, using Semantics based on W3C standards
- Manage the information lifecycle with tiered storage
- Build out applications in days and weeks that used to take months and years—without sacrificing data governance
- Lower costs by reducing time to deployment, as well as ongoing engineering, operations, and maintenance costs

Global enterprises use MarkLogic to re-imagine their data, using MarkLogic as a transactional and analytical database platform to run mission-critical apps at the heart of their data center. ■

MARKLOGIC

www.marklogic.com