

MarkLogic Mainframe Migration Accelerator

Mainframe Maintenance Is Hard – Migrating Away Was Harder

Getting off the mainframe is onerous. If it were easy, business would have done it decades ago. Until recently it was simply too difficult. In many cases, mainframes are built for the enterprise, scale massively, and hold data in hierarchical formats. Reworking mainframe data to fit a new relational database or living with an unreliable NoSQL system ranges from difficult to unacceptable.

The good news is these same **migration challenges** become **migration enablers** when moving directly from mainframe to MarkLogic®. MarkLogic skips the ill-fitting relational phase in between.

The value of our direct approach is compelling. MarkLogic customers realize significant value in key areas when they migrate to MarkLogic:

- Development of new business solutions is far faster and lower risk
- MarkLogic enables business transformation by removing legacy technology challenges and providing shared data services
- MarkLogic heals the “split brain” syndrome where data is split between historical mainframe databases and current operational systems by integrating your silos versus creating more silos
- Transferring to MarkLogic reduces costs in two ways: by developing faster, and eliminating mainframe maintenance fees faster
- Finally, MarkLogic creates data hubs that are **operational** to provide one “source of truth” that is actionable versus static, getting you closer to automaticity where it makes sense and better decision making in your business overall

Skip the Relational Intermediary

Often, the invisible blocker to getting off the mainframe has been the incompatibility between mainframes and newer relational databases. The formats are different. Relational technologies often cannot scale to mainframe loads.

Relational technology still has merit for well-behaved and perfectly constructed data fitting nicely in rows and columns. However, the era of complete relational database dominance is over for ill-suited tasks and ill-behaved data. It is simply inefficient and expensive. NoSQL, Hadoop, and related technologies — from document databases to graph stores, to data lakes — are firmly established as mainstream modern data alternatives.

By skipping the relational confines of rows and columns, mainframe applications now can migrate more easily — and with MarkLogic can retain their true enterprise qualities.



MarkLogic Is Uniquely Suited to Legacy Mainframe Migrations

MarkLogic is the only Enterprise NoSQL database well suited to mainframe transition:

- MarkLogic supports the same massive scale of the largest mainframe — and beyond. The MarkLogic platform does so by **scaling out** on commodity hardware, **not by scaling up on expensive custom hardware**.
- MarkLogic supports hierarchical, document data models, using JSON, XML, and RDF — all W3C standards that interoperate seamlessly with modern web architectures. MarkLogic also supports relational data for those mainframes that have them.
- Mainframes are built for the enterprise with ACID transactions, high availability, failover, backups, replication and decades of hardening behind them. We also do all that with Enhanced Security trusted in the most mission-critical systems for some of the most discerning clients in the world — from HealthCare.gov to trade stores at top 5 investment banks. All the features from ACID to HA/DR you need are in MarkLogic — and many new features you never dreamed of, like temporal versioning and integrated search with semantics, help get you operational in months versus years.

Unlike data lake technologies, MarkLogic is a true database with integrated search and an application platform, rather than a staging area or batch processing ecosystem. MarkLogic supports transactional query and update access, enterprise-class security, and HA/DR replication features to allow you to “run your business” on MarkLogic. MarkLogic works with Hadoop, but provides an operational feature set, rather than purely analytic.

Four-step Mainframe Migration Process

With our experience in mainframe migration, MarkLogic Consulting has evolved the following four-step process to migrate off a mainframe to MarkLogic Server. Our **Operational Data Hub** pattern and associated **Data Hub Framework** cover the software and design part of the migration. MarkLogic Consulting has the process and the know-how to make your migration faster and low-risk.

Determine the Overall Approach

There are three major directions to take when migrating off the mainframe to MarkLogic. The first step is to evaluate your needs.

1. Lift and Replace

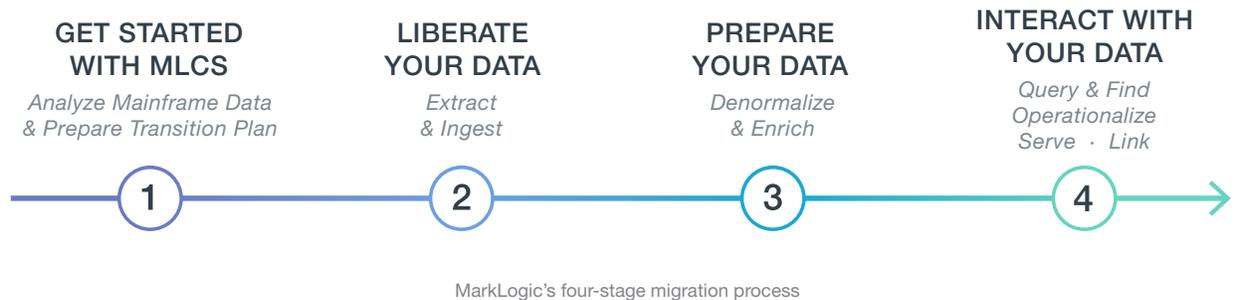
Sometimes, many inter-related applications live on the mainframe and are best migrated together. Or mainframe maintenance fees may be driving the migration, requiring a full migration to shut it down and save money. In these cases, a full simultaneous replacement is warranted.

2. Incremental Move

Specific functionality may be removed from a mainframe, leaving other functionality temporarily or permanently in place on the mainframe. If the application to migrate is coupled to the mainframe, some write-back capability must be added to the mainframe to accept updates and interact with the new applications. The approach is slower and supports the training and ramping of your internal developers to take over future stages by our expert consulting team.

3. Silo-busting

In addition to moving off the mainframe, you may combine the new application and reduce other data silos. Often, a legacy mainframe is too hard to modify, so a constellation of small systems has grown up around the mainframe, duplicating data and complicating operations. In this case, the small systems may be considered part-and-parcel of the same data migration job.



The Process

Once the overall approach is selected, MarkLogic Consulting will use the following process to understand, extract and operationalize your data in a new system.

Analyze

Together we analyze the data and the applications which use the data. Both identified data and applications will be moved to the MarkLogic platform so that the data is there, and can be used by the enterprise and users in the same way as before.

Extract and Ingest

Here is where the MarkLogic platform is truly unique and where the MarkLogic advantage begins to pay off. MarkLogic ingests data as *is*, including hierarchical VSAM or other files, or objects and entities joined from relational tables in a mainframe. What you don't see in this step is "transform." MarkLogic does pure ELT (extract and load — then transform later in ways that match your operations) to get data in fast, in any format, with **zero** modeling required up front.

If the mainframe has hard coding and relationships in code, these will be "liberated" by moving them into the data structures, where they belong for maximum benefit to the enterprise.

In the extract and ingest phase, data is often enriched with semantic markup to categorize and link incoming records (MarkLogic supports semantic standards such as RDF and OWL), and wrapped in an "envelope" that packages data and metadata together for easy use and querying.

Harmonize and Enrich

The imported *as is* models are grouped into sensible business entities and enriched to match business needs. Reference codes are typically merged into the records, and relationships expressed in a graph or RDF form. The result of this phase is a set of sensible business entities in the MarkLogic database, together with relations among them.

In particular, the approach here is in stark contrast to pure relational technologies, which scatter business entity groupings across many tables, with complex join logic to reconstitute a person, sales order, insurance policy, financial trade or other sensible entity. MarkLogic can actually match many mainframe approaches by storing grouped, hierarchical records. It's not the only way to store in MarkLogic, but makes mainframe migrations easier.

Query, Operationalize, Enrich, and Serve

MarkLogic shines here as it scales to mainframe loads, provides rich query capabilities, and makes the data operational. Simply put, it means users and applications can query the data or get exports that are often exactly the same as earlier mainframe data formats and services.



Enrichment may involve transforming data to new, harmonized forms, applying an **exacting security** design, data tagging to **add semantics**, entity linking to **uncover patterns**, semantics to uncover the **meaning in your data**, geo-tagging for map views, or simply getting **better and faster search results**. Collectively these activities make your data operational, allowing a move off the mainframe.

Data enrichment creates a more **robust foundation for effective operations** down the road by enhancing and improving existing data, to not only replace — but also enhance — previous mainframe operations.

Over a Decade of Experience

Whether your mainframe has 20 data sets or thousands, and whether information is hard-coded or clean, our experts can work with you to effectively extract your data, scaling to meet your needs. Moving your data from your mainframe into a more flexible database requires simple data extraction. Our experts will provide advice on the best approaches for accomplishing this based on your unique situation and end goals.

MarkLogic Consulting recently replaced an HR system with over 200 incoming and outgoing data feeds in a highly complex environment, using a pure MarkLogic solution. The project was successful, and we added temporal tracking along the way. The project was successfully delivered on time and on budget in **one fifth the time and cost of a relational approach**. MarkLogic consultants do this kind of complex transition every day, and have the background needed to reduce risk and deliver value in this complex area.

MarkLogic supports multiple implementation approaches, from a pure MarkLogic-led effort to a hybrid team of MarkLogic consultants integrated with your own staff and trusted implementers.

Security

It's not a separate phase. At MarkLogic, security is considered in every phase and is critical for modernizing and migrating data. MarkLogic uses role-based security for every record and every field so only authorized users see data, and only see what they are entitled to. Data is secured end-to-end: from encryption on disk, to security on every record and element, to HTTPS and other encryption on the wire. Any "upgrade" to a data platform with weaker security is arguably a downgrade instead.

Mainframe Transition Quick Start

The Mainframe Transition Quick Start service is an approach leveraging our quick start methodology that will get you up and running with the difficult task of transforming your legacy mainframe quickly. Our proven approach offers the most cost-effective, lowest-risk choice for customers looking for assistance in executing their mainframe transition with tangible and demonstrable results in months not years. Our service provides an approach for leveraging the MarkLogic platform to build a fully extensible data hub framework for business transformation. MarkLogic Consulting Services will engage side-by-side with your technical team to plan and execute your migration, including providing hands-on training, best practice advice, and validation of the implementation.

Many companies find any large technology change difficult. Migrating from decades-old, mature systems on mainframes is even more difficult. Fortunately, MarkLogic Consulting does exactly that all the time. Our consultants are expert at migrating data from silos, including mainframes, and bring a wealth of experience to the table.



About MarkLogic

MarkLogic is the world's best database for integrating data from silos. MarkLogic is an operational and transactional Enterprise NoSQL database platform, and organizations around the world rely on it to integrate their most critical data and build innovative applications on a 360-degree view of that data. 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. Headquartered in Silicon Valley, MarkLogic has offices throughout the U.S., Europe, Asia, and Australia.

Contact Us

To get started with MarkLogic Consulting Services or to get more information, you can contact your account representative, your local consulting leader or give us a call at [1-877-992-8885](tel:1-877-992-8885).

You can also email us at consulting@marklogic.com.

© 2018 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.

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