

PLAN FOR SUCCESS WITH HIGH-STAKES DATA PROJECTS

WHITE PAPER

Throw away your long-held traditions. The pursuit of your enterprise goals doesn't have to end as analysts predict—as one of many free falls into a growing pool of IT project failures. You can leverage all of your data to reach your objectives with less time and expense than you might imagine. The key is to start with the end in mind by undertaking your next high-stakes data project armed with the right technology and mindset to succeed. Change is good.



Contents

Executive Summary	1
Painful Truth: It's Time to Change Your Thinking About Data Integration Projects	2
Technology Hurdles	
Roadblocks of the Mind	
The Perfect Storm of Obstacles	
Decrease Risks by Rethinking Existing Technology Decisions	4
Adapt to the Right Mindset	
Adopt the Right Technology	
Stack the Deck in Favor of Your Enterprise	5
Implement A Federal Data Integration Project of Unprecedented Scope	
Build a Scalable Data Hub for the Delivery of Integrated Information in Real Time	
Rapidly Deploy a Unified and Secure Global Trade Store Within Six Months	
Set the Gold Standard for Dynamic Content Delivery at Scale	
Business Changes at the Speed of Thought – Be Ready to Change with It	12
More Information	

“ According to interviews and expert estimates, data scientists can spend up to 80 percent of their time mired in the mundane labor of collecting and preparing unruly data before it can be explored for useful nuggets.”

– Steve Lohr, New York Times

EXECUTIVE SUMMARY

The volume, variety and complexity of data is growing faster than ever before with the amount of information produced and consumed each day being nothing short of staggering. In fact, by the year 2020, about 1.7 megabytes of new information will be created every second for every human being on the planet.¹ Compounding the problems brought on by this state of global *infobesity*, the emergence of cloud computing, mobile platforms and the *Internet of Things* has left organizations hard-pressed to navigate effectively through a disorienting amalgam of messy data and continually shifting business requirements. Is it any wonder that 85 percent of companies fail to exploit their data to decisive competitive advantage?²

Risky by nature, many things can and do go wrong with the implementation of complex enterprise projects. *ComputerWorld* recently reported on a study of websites that cost them over \$10 million. The study showed that of the 3,555 projects across the public and private sectors, only 6.4 percent were successful. Over 50 percent of those projects experienced problems, such as going over budget, inability to meet deadlines or lack of user satisfaction.³

If the mere thought of planning an enterprise data integration project strikes fear into your heart, you are not alone. The stakes are high. According to Gartner, by 2018, 50 percent of the cost of implementing nearly

all new large systems will be spent on data integration.⁴ When it comes to implementing your own high-stakes data projects, are you heading toward certain disaster with the wrong technology and mindset?

If your IT projects find you relying on legacy technology and outdated thinking to build modern systems, chances are you're heading for a fall. In fact, another Gartner report found that through 2017, 60 percent of big data projects will fail to go beyond piloting and experimentation and will be abandoned.⁵ It's little wonder that many organizational behavior experts have reached the conclusion that fear of failure itself is a key contributing factor to the demise of many IT projects.

How will you manage your next mission-critical data integration project to ensure that it succeeds in achieving your business goals? How can you avoid having your next high-stakes data project become yet another statistic of Icarus-like crash and burn? Short answer: Pull together an implementation team with the *right mindset* and the *right technology* for success.

The good news is that by adopting both an outlook and technology that accommodates the reality of a business landscape continuously in flux, your organization can successfully implement data integration projects to reach your business goals.

This paper aims to reduce your risk of failure by sharing insights and pragmatic guidelines from industry

1 IDC: Digital Universe Study: Big Data, Bigger Digital Shadows and Biggest Growth in the Far East, December 2012 <http://www.whizpr.be/upload/medialab/21/company/Media_Presentation_2012_DigiUniverseFINAL1.pdf>

2 Gartner: Information Innovation: Innovation Key Initiative Overview, April 2012 <<https://www.gartner.com/doc/1991317/information-innovation-innovation-key-initiative>>

3 Computerworld: Healthcare.gov website 'didn't have a chance in hell,' October 2013 <http://www.computerworld.com.au/article/529631/healthcare_gov_website_didnt_chance_hell_/>

4 Gartner: Predicts 2013 Application Integration, 2013 <<https://www.gartner.com/doc/2238015/predicts--application-integration->>

5 Gartner: Advancing Business with Advanced Analytics <<https://www.gartner.com/user/registration/prospect?resId=3090420&srclId=1-6470978268>>

leaders who have learned hard lessons in their efforts to successfully deliver on their key complex data integration initiatives. The U.S. Department of Health and Human Services, Office of Inspector General, examined the development of Healthcare.gov in a detailed case study published in February 2016. You'll note that this paper draws upon this publicly available report to provide greater insight into the complexities faced and overcome during the implementation of the Healthcare.gov project.

PAINFUL TRUTH: IT'S TIME TO CHANGE YOUR THINKING ABOUT DATA INTEGRATION PROJECTS

Enterprises across the spectrum are working diligently to align their data assets with IT strategies for innovation and growth. The problem is that data never stops changing, growing and inundating your organization—making data integration across silos seem like a never-ending process. Just when you think you've finally gotten a handle on the requirements of your current project, more data comes flowing in, and more complexities are revealed. And although data architects work tirelessly to devise new ways to manage and leverage enterprise data, most organizations fail to exploit the data they've collected over the years to greater advantage.

TECHNOLOGY HURDLES

Barring the path to insight, the widespread deployment of purpose-built legacy systems to address specific needs has resulted in large data silos. Multiple copies of enterprise data spread out among these disparate silos have created challenges with data governance and accuracy, significantly inhibiting the ability for an organization to meet the needs of its clients in a cost-efficient manner.

Due to their complexity, enterprise data integration projects are risky, time-intensive and costly to implement. Organizations have attempted to lessen this complexity with multiple technology-based strategies, all of which have proven less than ideal. Included among them is the enterprise data warehouse




Figure 1. Complex data integration projects are high-risk

(EDW) that brings together data from disparate lines of business in support of downstream decision support analysis. EDWs often fall short because of the long tail required to move data into them from legacy operational systems. In addition, EDWs are typically built on rigid relational data structures that don't do well with changing source systems and downstream business requests.

Then there is enterprise application integration (EAI) and service-oriented architecture (SOA) for operational "run the business" business support. These frameworks are designed to enable the integration of systems and applications across an enterprise. Unfortunately, many applications can't communicate with one another in these types of systems, leading to identical data sets being stored over many locations. The end result of which can be messy, point-to-point connections that are difficult to manage.

Lastly, in an attempt to overcome these hurdles, some organizations are experimenting with data lakes (e.g., Hadoop) for large-scale storage of structured and unstructured data. The danger of data lakes is that they can quickly transform into quagmires as organizations load massive amounts of data into them without proper governance—leading to an unusable data swamp that can't be leveraged for operational value or business insights.⁶

⁶ [icrunchdatanews.com: How to Keep Your Data Lake from Becoming a Data Swamp, April 2016](https://icrunchdatanews.com/how-keep-your-data-lake-from-becoming-data-swamp/)
<<https://icrunchdatanews.com/how-keep-your-data-lake-from-becoming-data-swamp/>>

“ Half of IT projects with budgets of over \$15 million run 45 percent over budget, are seven percent behind schedule and deliver 56 percent less functionality than predicted.”

– McKinsey & Company

ROADBLOCKS OF THE MIND

Today we face the unfortunate reality that many of our IT projects will fail to meet deadlines and budget constraints. High costs and failed projects are the new norm, and according to McKinsey & Company, half of IT projects with budgets of over \$15 million:

- Run 45 percent over budget
- Are seven percent behind schedule
- Deliver 56 percent less functionality than predicted

Even worse, about 17 percent of IT projects fail so expensively that they can threaten the very existence of an organization.⁷

Why is it that these dismal stats go unchallenged as business as usual? One theory suggests that people at all levels of an organization are often averse to change when it pertains to the adoption of new technologies and processes. In fact, in an interview with Forbes, Andrea Simon, Ph.D., Cultural Anthropologist and CEO of Simon Associates Management Consultants, attests that it's difficult for organizational leaders and their employees to accept that “*limitations to growth are really self-imposed*” by the habits and practices that led to previous business successes. That may seem counterintuitive on the surface, but the crux is that “*past perceptions limit what they ‘see’ in the present.*”⁸

This tendency to look back can create a sense of blindness toward newer solutions that are better suited for the task. In the case of integrating modern data, enterprise architects tend to rely on the legacy

technology that worked for them in the past—in particular, relational databases—as the only way to create new systems. And even when they are willing to explore new alternatives, due to being so excited about using new technology, they stumble blindly into the future and overlook the importance of mature capabilities.

THE PERFECT STORM OF OBSTACLES

To illustrate, let's look at the State of Oregon's \$250 million failed attempt to develop its health insurance marketplace website, *Cover Oregon*, on relational database technology. Upon announcing closure of its website and subsequent plans to redirect state consumers to the fully functional federal health insurance marketplace website (HealthCare.gov), *Obamacarefacts.org* expounded with, “*The switch was recommended by an advisory panel after months of deliberation as fixing the website could cost up to \$78 million, while a switch to Healthcare.gov would only cost about \$7 million.*”⁹ According to Fortune, “*The state blames the software company's ‘incompetent work’ for the failure of its Cover Oregon healthcare enrollment website.*” Further, claiming that the state was unable to ever successfully launch this health insurance marketplace due to technical problems, the state alleges that the technology vendor committed fraud by convincing state officials to purchase “*hundreds of millions of dollars of [its] products and services that failed to perform as promised.*” The technology vendor responded with its own claims about the state's culpability in the data project's failure.¹⁰

7 McKinsey & Company: Delivering large-scale IT projects on time, on budget and on value, October 2012 <<http://www.mckinsey.com/business-functions/business-technology/our-insights/delivering-large-scale-it-projects-on-time-on-budget-and-on-value>>

8 Forbes: Why We're So Afraid of Change – and Why That Holds Businesses Back, April 2013 <<http://www.forbes.com/sites/womensmedia/2013/04/08/why-were-so-afraid-of-change-and-why-that-holds-businesses-back/#65d63d865cd0>>

9 ObamaCareFacts.com: Cover Oregon Marketplace Changes, April 25, 2015 Update <<http://obamacarefacts.com/insurance-exchange/oregon-health-insurance-exchange/>>

10 Fortune.com: Oregon sues Oracle over failed Obamacare website, August 2014 <<http://fortune.com/2014/08/22/oregon-oracle-spat-200-million-lawsuit/>>

“ We’ve all been there. If you always do what you’ve always done, you’ll always get what you’ve always gotten.”

– Tony Robbins

The result of this legal battle remains to be seen, but regardless of the outcome, with hundreds of millions of dollars in taxpayer money flushed down the drain on a critical data project that failed to see the light of day, there are a few burning questions that remain on the overall planning of the project: With a legally mandated program requiring the integration of a vast variety of data types, why was relational database technology automatically considered to be the best fit for the job? And what about the technology vendor, was it selected for the project based on its ability and willingness to align itself with the state’s goals and vision? Or, was the state’s choice in vendor driven more simply by familiarity with a technology that worked in the past? Questions that compel organizations to more closely examine why we do what we do can reveal surprising red flags of an organizational mindset that may not always work in alignment with your greater enterprise aspirations.

Every organization wants to leverage its data to innovate and outcompete in its market, but a fundamental fear of change and the inability to recognize and choose technology appropriate to the new realities your organization faces will stop even the best-laid plans in their tracks. Stacking the deck in your favor hinges on your ability to mitigate risks by ensuring technology and partnership decisions that advance your organizational goals and set the stage for your success.

DECREASE RISKS BY RETHINKING EXISTING TECHNOLOGY DECISIONS

ADAPT TO THE RIGHT MINDSET

Have you performed due diligence to ensure that you’ve chosen the right people to help your program succeed? The first step is to evaluate the mindset within your own organization and that of prospective partners to ensure you select the right people to implement your vision. The right people will be willing to readily adopt your enterprise’s spirit of innovation and balance it with your goals for today and tomorrow—all while mitigating risk. To increase the likelihood of success in implementing mission-critical data projects, it’s crucial to clearly communicate from the outset that a flexible, modern mindset among staff and partners is just as critical as adopting the right technology for the job.

ADOPT THE RIGHT TECHNOLOGY

Are you relying on yesterday’s technology to solve today’s problems? Flexibility and agility are key components to modern system design because they future proof the system and your investment. There is no way to predict the future, but the best way to compensate for the unknown is to create a flexible and agile system that can easily adapt to your changing requirements.

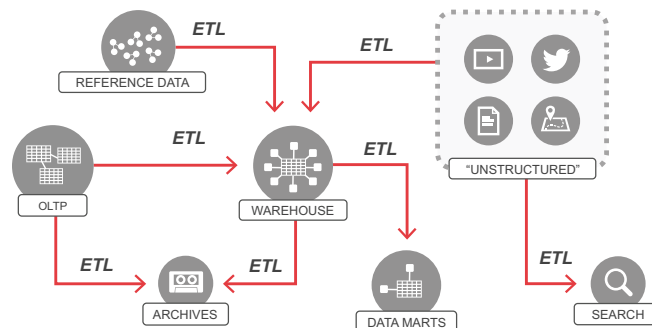


Figure 2. ETL is complex and time consuming

RDBMS

In use for over 30 years, relational database management systems (RDBMSs) won't disappear from the scene any time soon, but it's clear that a long history doesn't compensate for the inherent lack of agility within relational technology. The development of any new, agile system will require a flexible data management system.

Relational technology presents limitations for organizations needing a data platform that's capable of handling multiple sources of data, because they were created in an era when existing paper-based processes (e.g., payroll and ERP) needed to be automated. Process automation is no longer the goal. Today, organizations need the ability to combine multiple sources of data to create cloud-based mobile applications that can change quickly as needed. The rigid schemas that underlay RDBMS technology don't lend themselves to the fast-paced and constantly shifting needs of today's enterprise. Fixed data models combined with changing downstream systems create an instant incompatibility between the requirements of the business and the architecture of the project.

OPEN SOURCE

In an era of expensive IT projects that fail to launch, open source technologies can present an almost irresistible allure to cost-conscious decision-makers. In contrast to relational technology, the agile NoSQL (Not-only SQL) database market delivers more flexible data management and more cost-effective scaling for massive volumes of data. It encompasses a broad

field with dozens of entrants. The risk here is that these new open source vendor technologies were not designed to support mission-critical applications with the enterprise features (e.g., government-grade security, ACID transactions and HA/DR) required to support and safeguard your organization. Enterprise-level safeguards are vital because although consumers may not care if they lose track of a Tweet or an update to their Facebook status, they will be very alarmed to discover an incorrect bank balance or missing stock trade. So, before selecting an open source NoSQL provider, ask yourself how you will acquire these essential enterprise capabilities. Wiring open source technologies together results in higher development costs and is likely to yield a brittle, unreliable infrastructure that can't evolve alongside your changing business requirements over the long term. This is where open source gets tricky. The money you think you save on upfront licensing is really nothing more than a shifting of costs over to new and ongoing development.

STACK THE DECK IN FAVOR OF YOUR ENTERPRISE

In contrast to rigid RDBMS and risky open source NoSQL, the right database technology can eliminate data silos, bridge the gap between your enterprise analysis and operations and bring discovery to points of engagement. It will also handle multiple types of data, mixed workloads and scale on commodity infrastructure and the cloud to meet the requirements of today's enterprise—without business disruption.




Figure 3. Enterprise NoSQL harmonizes data across the enterprise

The missing piece of the puzzle, the MarkLogic® Enterprise NoSQL database, does all of this and ensures data integrity and reliability, enabling you to adopt a flexible approach to data projects that reduces the effort needed to reach your technical milestones.

A unifying platform for harmonizing data across your enterprise, flexible MarkLogic is designed to fit—without disruption—within your existing infrastructure. With this agile Enterprise NoSQL database, you can implement incrementally without the need to first identify every source of data and develop a common data model upfront. Built to enhance and simplify your data integration efforts, MarkLogic is *data-centric*, not function-specific, allowing one copy of data to serve many purposes.

Developed specifically for the enterprise, MarkLogic brings together the best capabilities of RDBMS and NoSQL to meet your organization’s enterprise requirements for government-grade security, ACID transactions, high availability, disaster recovery, agility, scalability and performance.

The best database platform for extracting value from *all* data—regardless of type, source or silo—MarkLogic answers your 21st-century data challenges with modern features built to solve them.

LEVERAGE LESSONS LEARNED FROM LEADERS

MarkLogic technology and expertise have enabled organizations across the globe to overcome some of the most complex data integration challenges the world over. Let’s review the issues and outcomes of a few recent high-stakes data projects implemented on this new generation database technology.

IMPLEMENT A FEDERAL DATA INTEGRATION PROJECT OF UNPRECEDENTED SCOPE

In stark contrast to the technology failure of the *Cover Oregon* health insurance marketplace, HealthCare.gov had a very rocky start, but ended its first operational year by providing health coverage to 7.1 million

Americans¹¹ in 14 weeks with a complete turnaround in just 10 weeks of going live. HealthCare.gov provides us with an unprecedented look at how the right technology choices in the beginning can be a huge asset when an extremely large project goes horribly wrong.

THE CUSTOMER

Centers for Medicare and Medicaid Services (CMS)

THE SITUATION

With an unprecedented scope, new and evolving policies to comply with, escalating costs and a mandated near-impossible deadline, the stakes were high. Henry Chao, CMS Deputy CIO and Director, Office of Information Services, and technology lead on the initial project rollout, remembers the intense pressure:

“Imagine the following IT project challenge—create a system scalable for tens of thousands of concurrent active users with 99.9% uptime that touches and integrates with at least 20 other existing massive systems that are not under your control. You must execute delivery of this system within a culture where hundreds, maybe even thousands, of stakeholders provide ongoing input, and not one of them has ultimate decision-making power. Because of this culture, the requirements for the system keep changing as you build it. You must also deliver this complete system by an unchangeable date, and everyone is watching for you to fail.”¹²

ENTERPRISE OBJECTIVE

Tasked by the federal government to launch a nationwide health insurance marketplace mandated by the Affordable Care Act, CMS had to bring together data from many states, numerous federal agencies and insurance companies to create a seamless online shopping experience designed to extend access to healthcare for millions of Americans.

THE CHALLENGE

CMS relies on major prime contractors to build, manage and maintain all of its systems. CMS additionally has many highly knowledgeable employees and leaders that have been with the organization for years and have

11 Washington Post: At least 7.1M signed up for 2015 Obamacare plans so far, December 2014
<<https://www.washingtonpost.com/news/wonk/wp/2014/12/30/at-least-7-1m-signed-up-for-2015-obamacare-plans-so-far/>>

12 InformationWeek: Healthcare.gov: Hard-Earned Lessons for CIOs, March 2016
<http://www.informationweek.com/strategic-cio/healthcaregov-hard-earned-lessons-for-cios/d/d-id/1324456?page_number=1>

“ HealthCare.gov is not just a website. That's really an oversimplification. HealthCare.gov is a complex eligibility verification and determination application that is integrated with a data services hub that serves as a broker for all the requests and responses coming from authoritative sources.”

– Henry Chao, *InformationWeek*

successfully deployed massive systems that touch every American. The stakes are extremely high for every system the organization is mandated to deploy, and every leader and employee is deeply aware of their responsibility to the American people. A key part of this responsibility is balancing risk.

A reasonable “go to” strategy to balance risk is to use proven technologies and processes that have worked in the past. The prime contractors who understand this cultural aspect of the organization bid on these contracts with traditional technologies, including RDBMSs for handling all of the disparate data that

would need to be managed by the Marketplaces and the Data Services Hub (DSH). CMS, on the other hand, knew what lay ahead and knew that relational technology and traditional development processes would not work. First, the relational technology couldn't handle the wide variety of data formats flowing into the system from states, insurers and other federal agencies. Then, it also became clear that this database system wouldn't be able to scale up quickly once consumers hit the site in droves on the first day of open enrollment. Taken as a whole, these shortcomings meant that applications built on top of the existing relational database wouldn't function in a way that

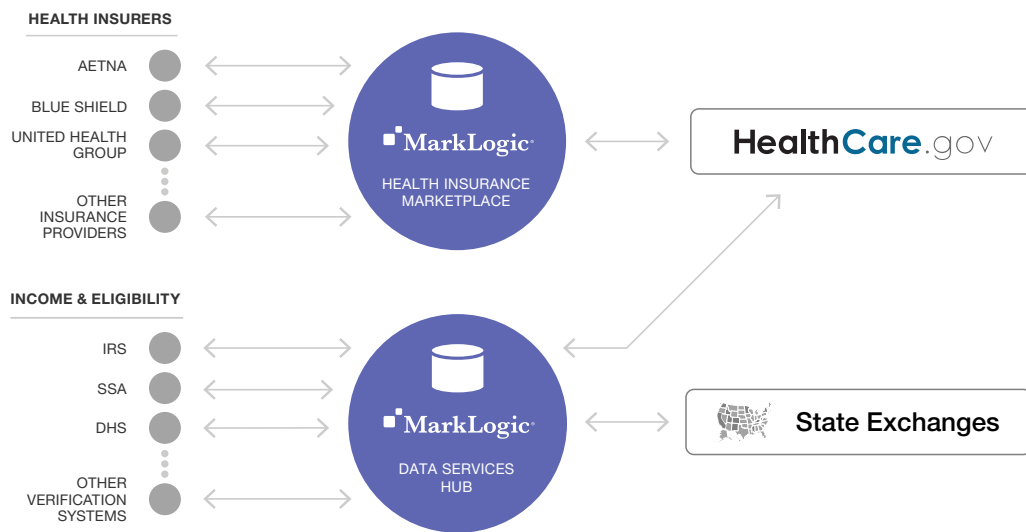


Figure 4. A technology platform capable of enrolling millions of users in new healthcare plans

13 U.S. Department of Health and Human Services: Healthcare.gov Case Study of CMS Management of Federal Marketplace, February 2016 <<http://oig.hhs.gov/oei/reports/oei-06-14-00350.asp>>

14 U.S. Department of Health and Human Services: Healthcare Gov. Case Study of CMS Management of Federal Marketplace, February 2016 <<http://Oig.Hhs.Gov/Oei/Reports/Oei-06-14-00350.Asp>>

adequately supported CMS's mission to meet the legal requirements of the Affordable Care Act.

CMS and Henry Chao, its Deputy CIO, had no option but to award contracts to these prime contractors and then provide them with alternate technical direction. They specified MarkLogic as the database platform in a January 2012 technical direction letter.

The U.S. Department of Health and Human Services, Office of Inspector General's case study on HealthCare.gov supports this decision, citing, "The key benefit of the nontraditional NoSQL platform is its potentially greater capability, in that it can allow more data to be transferred at a time and can be easily expanded to include more data or users."¹³

Replacing the RDBMS that couldn't meet CMS's requirements, MarkLogic was selected because it can deliver the flexibility and performance required for implementation of a complex project on this scale. But, being the newer technology on the block among established government vendors was a challenge. In fact, lacking experience with NoSQL database technology, one prime contractor misperceived MarkLogic as a potential risk to timely implementation, reporting that "it would be unable to fully replace staff expertise on traditional databases with equivalent expertise with MarkLogic's NoSQL platform" and recommending that CMS set up a contingency plan.¹⁴

CMS proceeded with MarkLogic as planned. And ultimately, other key contractors on the project who were more receptive to newer database technology were able to quickly leverage MarkLogic for the successful implementation of the Data Services Hub and the launch of HealthCare.gov.

The experience of CMS raises some thought-provoking questions regarding risks brought on by fear of change. In this case, one of CMS's contractors perceived a risk in adopting newer technology. Fortunately, Henry Chao made the decision to break with the norm, understanding the greater risk of failing to meet his organization's requirements and deadlines; and that

legacy technology would not work. This proved wise:

- The flexible MarkLogic platform was able to handle the complexity of integrating disparate data for both the Marketplace and the DSH, despite changing data sources and requirements.
- Even in the face of network data center failure, the MarkLogic database functioned flawlessly—safeguarding CMS with zero data loss.
- The DSH—whose development was led by one of the contractors who was more receptive to new technologies—has run seamlessly since launch.
- Despite the early failure of the consumer-facing Marketplace application, CMS was able to rapidly turn the system around. According to Henry Chao, "When things were bad, we had the option to pivot, to scale out of a poorly written application." MarkLogic "gave us a set of options that would not have been possible with other technologies."¹⁵

KEY RESULTS

According to CMS, the 2016 health insurance open enrollment period ended with over 12.7 million consumers enrolled in marketplace coverage through the HealthCare.gov website.¹⁶ With MarkLogic, CMS was able to:

- Meet the congressional deadline for the delivery of the nationwide health insurance marketplace mandated by the Affordable Care Act.
- Launch a highly scalable operational data hub with 100 percent data integrity and the stability to support 280,000 concurrent users.
- Securely integrate a wide variety of structured and unstructured data from many systems into a seamless user experience capable of handling 6,500 transactions per second.
- Extend access to health insurance coverage for millions of Americans.

¹⁵ MarkLogic: The Untold Story of Rescuing Healthcare.gov, December 2015
<http://www.marklogic.com/resources/the-untold-story-of-rescuing-healthcare-gov/resource_download/whitepapers/>

¹⁶ CMS.gov. Health Insurance Marketplace Open Enrollment Snapshot – Week 13, February 2016
<<https://www.cms.gov/newsroom/mediareleasedatabase/Fact-sheets/2016-Fact-sheets-items/2016-02-04.html>>

BUILD A SCALABLE DATA HUB FOR THE DELIVERY OF INTEGRATED INFORMATION IN REAL TIME

THE CUSTOMER

Aetna

THE SITUATION

Aetna is a large health insurer that provides coverage for over 23 million people through over one million providers. Not only does this mean that it processes billions of claims every year, but it also has to manage human resources data for 35,000 of its own employees. The 2015 acquisition of Humana added yet another source of significant new data and requirements.

ENTERPRISE OBJECTIVE

Aetna's IT organization was under increasing pressure from the business to integrate all of its human resources data (e.g., employee data, payroll data and many other systems) into a central place and have it delivered to downstream systems for real-time apps and batch analytics. Its objective was to build a scalable human resources data hub for the delivery of integrated human resources information in real time.

THE CHALLENGE

Much of Aetna's existing data was stored in traditional relational database models. Integrating this data would typically require the development of complex point-

to-point data exchange solutions involving highly complex ETL processes that all had to be written upfront. The approach would not only be costly and time-consuming, but the rigid constraints of relational databases would ultimately result in a brittle system that would require a great deal of rework to implement any changes in the future.

KEY RESULTS

With MarkLogic, Aetna was able to:

- Successfully integrate over 24 different HR data sources composed of about 140 separate data feeds on a system that handles a few terabytes of data and has a throughput of about 50 GB of data per day.
- Make use of a flexible data model it could develop in an agile manner to overcome the challenges of complex ETL processes.
- Meet implementation costs and timeline constraints without sacrificing crucial enterprise requirements, such as ACID transactions security, auditability and transparency.
- Easily integrate new data sources and adapt to changes as needed on its new future-proof database platform.
- Bring increased satisfaction to customers at a lower cost.

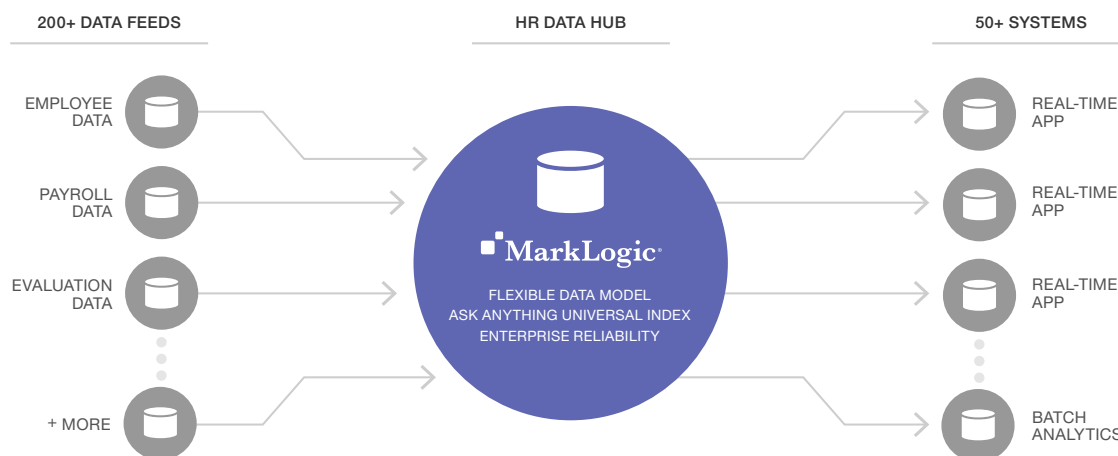


Figure 5. Integrated HR data, real-time delivery at scale

RAPIDLY DEPLOY A UNIFIED AND SECURE GLOBAL TRADE STORE WITHIN SIX MONTHS

THE CUSTOMER

Deutsche Bank

THE SITUATION

With more than 100,000 employees in over 70 countries, Deutsche Bank is a global banking and financial services company headquartered in Germany. The banking leader offers financial products and services for corporate and institutional clients along with private and business clients.

ENTERPRISE OBJECTIVE

To attain a single source of truth across its multiple trade types, Deutsche Bank needed to integrate its data and create an operational trade repository (global trade store) for its integrated post-trade processing. The bank's downstream business lines would access the unified data to fulfill a number of reporting objectives, including regulatory compliance requirements.

THE CHALLENGE

Deutsche Bank has many trade-related data inputs that are stovepiped into multiple lines of business, all of which use the data in a number of different ways. The bank's challenge was to harmonize its data, create a single version of the truth and ensure data consistency.

Deutsche Bank tried to achieve these accomplishments by using RDBMS technology, but the project failed due to the inability to agree upon a required common schema and shifting business requirements. A more flexible data integration technology was needed.

KEY RESULTS

With MarkLogic, Deutsche Bank was able to:

- Successfully integrate over 30 trading systems into a flexible and unified global trade store—achieving a single source of truth.
- Reduce costs with an agile trade store that easily expands its provisioning capacity as needed while eliminating the need for massive ETL efforts that would have been required under its previous relational technology.
- Run the bank on a single, enterprise-ready database platform that ensures secure and consistent transactions.
- Save time by quickly integrating all types of structured and unstructured data and then rapidly deploying the initial system within six short months.

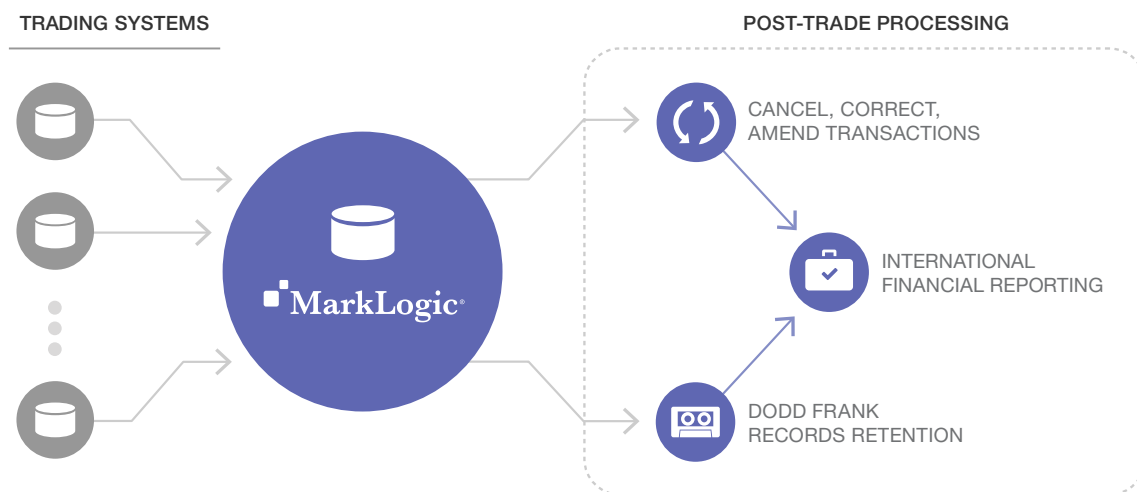


Figure 6. Integrated post-trade processing

SET THE GOLD STANDARD FOR DYNAMIC CONTENT DELIVERY AT SCALE

THE CUSTOMER

British Broadcasting Company (BBC)

THE SITUATION

The BBC is the world's leading public service broadcaster, operating 10 national and 40 local radio stations, digital applications and an extensive website.

ENTERPRISE OBJECTIVE

The BBC wanted its online coverage of the 2012 Summer Olympic Games to include an athlete information website designed to deliver an unparalleled level of personalized interactive content to any user device in real time. Having previously relied upon a RDBMS to run its static publishing website, the BBC understood that its rigid database platform would not allow it to effectively manage and deliver the steady stream of incoming and outgoing variable data required to create the enhanced interactive user experience it sought to provide for Olympic viewers and athlete fans.

THE CHALLENGE

To attain its goal of delivering a full-featured, interactive website, the BBC needed to develop an agile database platform scalable enough to support tens of thousands of transactions per second. The BBC's efforts were curtailed by a wide range of roadblocks, including a large amount of complex data stored in disparate systems, and an outdated IT infrastructure incapable

of delivering dynamic, multi-format content (e.g., video, editorial and statistical data) to audiences. In addition, there was a lack of database technology that was agile and powerful enough to accommodate millions of website users with anytime, anywhere access to all of its Olympic-athlete information and coverage.

KEY RESULTS

With MarkLogic, the BBC was able to:

- Speed time to production by reducing the build time and infrastructure constraints required by its existing relational technology.
- Leverage ACID transactions to provide stable, uninterrupted digital coverage with no lost data, degradation in site performance or downtime.
- Expertly handle over 25,000 transactions per second and 45 billion requests to serve up 10,000 dynamic pages to website visitors.
- Create a highly scalable and personalized digital user experience that made it possible to deliver an unprecedented amount of interactive content (2.8 petabytes delivered on peak day) to millions of users over the device of their choice.

Having had success with MarkLogic in the deployment of its dynamic delivery platform, today BBC has extended its use of MarkLogic by leveraging it to improve the performance of its popular iPlayer video on-demand service.¹⁷



¹⁷ Computerworld UK: BBC iPlayer sees performance uplift after relaunching on NoSQL database, June 2104
<<http://www.computerworlduk.com/news/data/bbc-iplayer-sees-performance-uplift-after-relaunching-on-nosql-database-3524893/>>

BUSINESS CHANGES AT THE SPEED OF THOUGHT – BE READY TO CHANGE WITH IT

Many things *can* and *do* go wrong with the implementation of complex enterprise projects—but your technology doesn't have to be one of them. Don't settle for failure as the statistical norm. If chosen wisely, you can leverage your database as a decisive means for mitigating risks and bringing your company closer to its aspirations. The key is to prepare your project to win by pulling together the *right* people to build upon the *right* technology.

With business changing at the speed of thought, success hinges on your ability to deliver value more quickly than your competitors. To do that, build your critical projects upon a solid but agile data platform that will work for you today and evolve with you into the future. Don't get caught up in a false sense of security by attempting to use legacy technology that has worked in the past to build a system around today's modern data; it wasn't built to do the job. Instead, by adopting an outlook and technology that accommodates—rather than ignores—the quickly and constantly changing business landscape, your organization will execute successful data projects that empower you to attain your business goals.

MarkLogic is an operational and transactional Enterprise NoSQL database that has enabled organizations in every industry to integrate data from silos more quickly and cost-efficiently than ever before. We help solve the world's most complex data challenges by equipping business leaders with a unified, 360-degree view of their data, enterprise readiness, agility and scalability.

Cast off legacy mindsets and technology traditions that no longer take you where you need to go. With the right team and technology developed specifically for today's modern data requirements, not only will your

high-stakes projects succeed, you'll leverage all of your enterprise data to greater competitive advantage. That's a welcomed change.

MORE INFORMATION

Relational databases were not built to handle the challenges presented by the massive amounts of disparate and intricate data we use and produce each day. But, you can leverage all of your data to reach your enterprise goals with less time and expense than you once thought. Learn how MarkLogic's new generation database technology is clearing your path to successful data projects:

Beyond Relational

With RDBMS, implementing even a small change could morph into a costly nightmare. That's why leading organizations are moving beyond relational in favor of a flexible database technology designed to reduce data modeling – and potentially save your business millions. Read the white paper:

<http://www.marklogic.com/resources/beyond-relational/>

Enterprise NoSQL for Dummies

Learn how databases incorporate semantic technology to make it possible to solve big data challenges that traditional databases aren't equipped to solve.

<http://info.marklogic.com/nosql-for-dummies.html>

Solve complex data challenges

A new generation technology designed for today's modern data types, MarkLogic is the only Enterprise NoSQL database. Built with a flexible data model to store, manage, and search all of your data, without sacrificing any of the data resiliency and consistency features of RDBMS. Visit us on the web to learn more:

<http://www.marklogic.com/what-is-marklogic/>

© 2016 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