CHECKLIST 2022

Best Practices for a Modern Data Layer in Financial Services

By David Loshin

redis

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

Best Practices for a Modern Data Layer in Financial Services

By David Loshin

Organizations solely relying on legacy systems, especially in the financial services industry, are facing increasing challenges affecting current operations and competitiveness. Limitations in accessing and integrating data across the systems supporting different lines of business due to system complexity and data latency impact obligatory reporting and diminish the ability to address customer-facing needs (such as real-time personalization and recommendations, and fraud monitoring, prevention, and alerts). Because these limitations ultimately contribute to market erosion or missed business opportunities, organizations are increasingly considering modernizing their legacy transaction processing systems.

Digital-savvy millennials, whose banking experiences are limited to virtual interactions via smart devices, expect full integration among an array of different applications and interaction patterns with their financial service providers.

Six steps for data layer modernization:

1. Identify data accessibility gaps impeding competitiveness
2. Acknowledge the need for a consistent omnichannel customer experience
3. Envision information models for real-time business processes
4. Consider ways to accelerate adoption of modern computing services
5. Embrace data platforms supporting low-latency data accessibility
6. Incorporate an operational data layer designed for financial services organizations

However, organizations acutely dependent on legacy, on-premises mainframe platforms will discover that aging technology increasingly impedes their response to newer financial technology companies, popular payment solutions (e.g., Apple Pay or Google Pay) or digital wallets, new digital currencies (such as Facebook's Libra), and companies forming partnerships with banks while establishing their own financial services options.

Remaining bound to legacy systems limits these organizations to batch processing. They are unable to dynamically scale performance when necessary, and must segregate transaction processing from the integrated analytics needed to drive new business opportunities. Modernization in the financial services industry is not just about accelerating transaction performance; it is about providing a streamlined pathway for digital transformation that integrates with multiple third-party systems and APIs and supports concurrent transaction and analytical processing—all executed in real time.

At the same time, financial services companies can balance organizational risk factors through a measured approach to modernization. This approach embraces on-premises *and* cloud platform architectures providing high availability, automated failover, and rapid data access across a hybrid (combined on-premises *plus* cloud) deployment. The approach also sets the stage for adopting a microservices architecture while providing fully integrated security and protection.

This checklist report offers steps for preparing for data-layer modernization that is independent of underlying platform choices. The checklist recommends ways to develop a strategy for transitional modernization—without disruption—that is consistent with the organization's risk mitigation profile while addressing the demand for real-time customer experience and other critical business functions.

1 Identify data accessibility gaps impeding competitiveness

Within the past decade, the consumer-facing facet of financial services has been radically transformed. In the past, the predominant interface between the customer and the financial institution was a brick-and-mortar bank branch where humans manually processed bundled transactions while providing a paper trail. With a focus on speeding customers through the pipeline, cross-interaction among services was limited.

Today's start-up financial technology (FinTech) companies provide virtual services, enabling real-time transaction processing via mobile apps that provide nearly continuous connectivity between customer and financial institution. Straight-through processing speeds automated transactions; business rules that are fully integrated across products and services trigger personalized solicitations for customer engagement. Big tech companies are integrating payment methods and incentivizing consumers to adopt their credit products.

Older financial institutions that haven't begun their digital transformation will soon find themselves unable to contend effectively in an increasingly competitive market. One key barrier to modernization is the lack of awareness of how platform, system, or vendor choices often become synonymous with the business processes themselves.

An example is a process that batches transactions into a file, uploads the file to a centralized server, and then schedules the file's bundled transactions for processing. Although there is no *intrinsic* need for the transactions to be processed in scheduled uploaded batches, the perceived tight coupling of the implementation with the business process itself hides this technical dependency.

The first step towards establishing a modern data layer involves assessing the current legacy environment to identify and eliminate latent systemic constraints. Create an inventory of your most critical business processes and review each corresponding application's automated and manual information touchpoints. Ask whether any existing technology choices or legacy system constraints associated with the availability, accessibility, latency in delivery, or the integrability of data impede the enterprise from improving its real-time customer experience or speeding its risk mitigation processes such as real-time fraud prevention. Prioritizing opportunities and scheduling projects that eliminate such technical dependencies by modernizing the data layer will position your organization to more effectively compete with newer financial services organizations.

2 Acknowledge the need for a consistent omnichannel customer experience

In almost all industries that deal with retail clients, customers have become accustomed to immediate (and in some cases, continual) engagement with their product or service providers using multiple devices, channels, and access methods.

Consider retail banking clients desiring visibility into the array of fully integrated account services. One customer might access her checking, mortgage, and credit card accounts to arrange for automated transfers and payments, while another might want to tightly couple money market accounts with security trading accounts to rapidly fulfill trades. Yet whether customers interact via telephone, email, an online e-commerce platform, chat, or an in-person branch visit, they find comfort when presented with a consistent omnichannel experience.

The customer has four fundamental expectations associated with that experience:

- Accurate and up-to-date information
- Fast response time
- Access to the full array of products and services
- Offerings that are appropriately personalized

A financial institution unable to accommodate these expectations may expect lower profitability as customers drift to more electronically agile companies such as FinTech disruptors/digital-only neobanks (for example, Chime, Nubank, or Varo), service providers for trading (including Robinhood or Acorns), or mobile payment services companies (such as Venmo).

Supporting a real-time, consistent omnichannel experience implies that no matter the interaction scenario, the financial institution can quickly access a view of the customer's data integrated from across different areas of the business.

This view must consistently reflect a coherent, up-to-date perspective of all customer accounts. The view's quality must be impeccable because customers will not tolerate inaccuracies. It also must simultaneously present accurate personalizations and recommendations.

A modern data layer speeds data access and employs in-memory management of data accumulated from across a network of systems. This enables rapid access to an integrated view of the customer's account landscape that facilitates account review but simultaneously supports embedded integrated analytics driving personalized suggestions and offers.

3 Envision information models for real-time business processes

Once you have started envisioning the ability to access customer data and customer profiles that can be managed and updated in real time, you can rethink other real-time business processes that share two key characteristics:

- They depend on (but currently don't support) rapid access to fully visible information models
- Their operational needs align with back-office processing of transactions impacting the actual system of record

There are multiple examples for identifying and taking advantage of discoverable business opportunities throughout the customer experience: making personalized offers, targeting advertising, creating discounted product/service bundles, dynamically pricing or configuring products and services, as well as offering special or limited time promotions. These information models encapsulate integrated customer views, configurations for materialized data models (which may employ traditional RDBMS and NoSQL-based databases), embedded stream processing models, or fast search; they may include repositories collecting the features necessary for machine learning and AI algorithms. Supporting a variety of information models simplifies the analytics life cycle by enabling data consumers to quickly find managed data assets and rapidly access and stream them on the same platform.

Real-time processing can meet additional immediate business needs. One such domain is fraud detection—enabling automated fraud analysis and prevention, which becomes increasingly complex with a growing network of partner institutions, distributed customers, and number of integrated applications. Another domain is compliance, such as real-time monitoring for meeting the requirements of know your customer (KYC), anti-money laundering (AML), or foreign account tax compliance (FATCA). A third domain is the emerging open banking regulations intended to make it easier for sharing customer data among different providers and third-party tools. Data sharing through open banking APIs creates opportunities for promoting new products and services.

In a legacy environment, supporting these business processes is cumbersome. Producing the integrated view is challenging when the required data must be sourced from a variety of systems. The need to fold structured and unstructured data together adds to the challenges, particularly when some of the processes are implemented manually.

A modern data layer reduces the overall complexity of these processes by employing representations of customer information customized for real-time operations. A modern data platform maintains a consistent view of data pulled from a variety of sources while providing real-time data integration, management, and access. This improves productivity using continuous pattern analysis and execution of integrated analytics for pattern analysis, personalization, and recommendations, as well as real-time execution of compliance and fraud analysis applications.

4 Consider ways to accelerate adoption of modern computing services

An institution's desire to renovate its operational platform and enable real-time operations must take into account how a modernized environment can shorten response times and process execution while remaining cost-effective. Because it promises greater scalability and computational flexibility, cloud computing has become an attractive platform for an organization's digital transformation. Cloud computing essentially provides unlimited scalability for computing and storage resources, and effective management guardrails can help you manage costs.

That being said, although the ultimate destination may be the cloud, not every organization is prepared to migrate its entire application landscape at once. Although adopting cloud services can provide some benefits, risk-averse organizations will think twice about an "all or nothing" cloud migration. Instead, they will consider a hybrid model that blends traditional on-premises platforms as the systems of record, managed private cloud computing and storage environments (to take advantage of increased capacity and performance), and public cloud implementations (when demand for computing resources spikes).

This suggests considering a transitional strategy of incremental adoption of cloud services that can add value immediately within a hybrid environment. This strategy includes adopting application architecture components that accommodate the hybrid approach and eases the path to a future cloud deployment, such as with microservices and containers. Microservices can reduce the complexity of application development by breaking down each function into an independent service. In turn, these microservices can be deployed in a variety of architectural paradigms, ranging from bare metal and virtual environments to containers. Containers are a lightweight method for packaging applications independently of the underlying hardware and software stack.

Microservices and containers allow developers to modularize their applications so they are easily migrated across and among platforms. Transitioning to a microservices approach provides a runway for cloud adoption, especially if you need different data models/structures for different microservices. The result is a more flexible application architecture that is deployable across a hybrid environment and is ripe for alignment with a modern data layer.

5 Embrace data platforms supporting low-latency data accessibility

Our prior discussion has consistently maintained the critical importance of low-latency data access. We have suggested that even with the highest-performing computing resources, data latency will bog down the process and reduce an enterprise's ability to react and respond to events and opportunities, and to execute business critical processes in real time.

The key to real-time operation is the data layer, which can accelerate data access, employ multiple models layered on top of existing data sources (and their corresponding source models), and take advantage of intelligent state information management to help in speeding data access. For example, consider techniques for predictive access of customer data (driven by state information associated with user identity, IP address, device IDs, and so on) providing the ability to effectively prefetch the data to populate a conformed customer profile that is ready to respond to customer actions.

An in-memory database system or data layer will ensure low-latency data access. Yet not all in-memory data layers are the same. Look for an in-memory data layer that supports applications' needs for low-latency data accessibility that leverages intelligent memory management and application caching to support performance and scalability in a hybrid, distributed environment. Specifically, consider tools with the following characteristics:

- **High and scalable performance.** Maintaining data in memory reduces data access latency and speeds performance and throughput.
- **Persistence.** Simplistic in-memory stores do not necessarily provide persistence, which can become a disadvantage. Seek data layers that support persistence without sacrificing performance.
- **High availability.** A perceived risk with an in-memory store is its ephemerality, which suggests low fault tolerance. Seek a data layer that combines in-memory performance that can recover from failures.
- **Adaptability of the memory hierarchy.** Clearly, as data volumes grow, the amount of information will exceed the system's available memory. Look for a product that can tier data along the memory hierarchy, leveraging DRAM, solid-state disk storage, persistent disk storage, and more, according to the frequency of data access.
- **Flexibility in platform deployment.** As suggested in the last section, adoption of a microservices approach coupled with containerization can streamline application migration. Look for an in-memory solution that can be deployed in a variety of platforms and containers.

6 Incorporate an operational data layer designed for financial services organizations

For financial institutions with a history of dependence on legacy platforms, opting for the cloud is only one part of a modernization strategy. Risk-averse organizations recognize the difficulty of overcoming the hurdles of modernizing legacy

systems by replicating their operational systems directly in the cloud. They realize that modernization can be accomplished through incrementally transitioning the application architecture to ease any migration. This includes the adoption of a microservices approach that decouples the application from the platform and enables deployment across an array of platforms using containers.

This suggests a tactical modernization approach that targets value-added applications that leverage analytics coupled with integrated data to support real-time business needs. A microservices approach allows your development teams to begin with specific use cases that drive services design, and enables expansion to new applications in a controlled yet risk-averse manner.

Examples include shifting towards customer-focused digital banking to meet the needs of digital-savvy clients, real-time compliance monitoring and fraud applications, as well as ensuring consistency in real-time reporting. These applications initially take advantage of rapid access to a selected set of customer records, often offloaded from the primary databases or systems of record and cached within an operational data layer. The accuracy and precision of analytics models can be increased using additional microservices-based applications using the same data such as machine learning/AI algorithms.

In essence, consider how an operational data layer coupled with microservices can be reused to support multiple business processes and workflows. This operational data layer supports rapid, low-latency accessibility to operational or transactional data without the overhead associated with a fully ACID-compliant relational database. This facilitates the introduction of new applications to support business imperatives without disrupting existing backend operations.

Afterword

Risk-averse organizations are faced with a challenge—modernize to address the specter of becoming less competitive but do so in a non-disruptive manner. Organizations such as those in the financial industry must transition their foundational architecture in a way that eases the adoption of newer techniques to positively influence the customer experience, enable real-time analytics for mitigating risk and ensuring compliance, and support the mandated requirements for statutory and regulatory reporting.

Identifying technology dependencies that impede modernization is a prelude to developing a modernization strategy. At the same time, recognize that maintaining competitiveness with emerging technology-based institutions depends on an application development framework that can access the full array of systems of record containing the customer's products and services, ensure data quality and accuracy, and quickly produce analytics results for personalization, recommendations, and compliance.

Accelerate adoption of modern computing techniques by considering a microservices architecture layered on top of an operational data layer. Partner with technology vendors that have developed an in-memory data management layer designed to interoperate with a microservices architecture to support tactical modernization. A high-speed, low-latency operational data layer enables high-performance computing and advanced

integrated analytics. This data layer supports real-time customer interactions, recommendations, personalization, and other key business processes such as compliance and fraud detection and prevention. At the same time, look for a data layer providing the high availability, fault tolerance, and enterprise-grade security that are core expectations for risk-avoidance.

About our sponsor

redis.com

Businesses are more digital today than ever before. They need to build, deploy, and run real-time services in order to stay ahead of the curve. The notion of real-time isn't just a nice-to-have anymore. It's an expectation. It's what sets a merely good user experience apart from a great one. A real-time data layer is a critical enabler in creating those real-time experiences.

Redis makes apps faster by creating a data layer for a real-time world. We're the driving force behind Open-Source Redis, a much-loved, in-memory database, and the commercial provider of Redis Enterprise, a real-time data platform.

Redis Enterprise powers real-time services for over 8,500 organizations globally. It builds upon the simplicity and speed of Open-Source Redis along with an enterprise-grade data platform that offers robustness of modern data models, management, automation, performance, and resiliency to deploy and run modern applications at any scale from anywhere on the planet.

About the author

David Loshin, president of Knowledge Integrity, Inc., (www.knowledge-integrity.com), is a recognized thought leader and expert consultant in the areas of data management and business intelligence. David is a prolific author regarding business intelligence best practices as the author of numerous books and papers on data management, including *Big Data Analytics: From Strategic Planning to Enterprise Integration with Tools, Techniques, NoSQL, and Graph* and *The Practitioner's Guide to Data Quality Improvement*, with additional content provided at www.dataqualitybook.com. David is a frequently invited speaker at conferences, web seminars, and sponsored websites and channels. David is also the Program Director for the Master of Information Management program at the University of Maryland's College of Information Studies.

David can be reached at loshin@knowledge-integrity.com.

About TDWI Research

TDWI Research provides industry-leading research and advice for data and analytics professionals worldwide. TDWI Research focuses on modern data management, analytics, and data science approaches and teams up with industry thought leaders and practitioners to deliver both broad and deep understanding of business and technical challenges surrounding the deployment and use of data and analytics. TDWI Research offers in-depth research reports, commentary, assessments, inquiry services, and topical conferences as well as strategic planning services to user and vendor organizations.

About TDWI Checklist Reports

TDWI Checklist Reports provide an overview of success factors for a specific project in business intelligence, data warehousing, analytics, or a related data management discipline. Companies may use this overview to get organized before beginning a project or to identify goals and areas of improvement for current projects.

tdwi
Transforming Data With Intelligence™

A Division of 1105 Media
6300 Canoga Avenue, Suite 1150
Woodland Hills, CA 91367

E info@tdwi.org

tdwi.org