



Unlocking banking's data advantage

From isolated islands to
connected, reusable assets



Introduction

Digital technologies have played a transformative role in financial services growth over the past thirty years, but data is the sector's digital Achilles heel. The inherent complexity, inefficiency, and risk associated with managing data have limited potential digital productivity gains by an estimated 40% to 50%.

The artificial intelligence ("AI") banking era only intensifies the issue. Customers increasingly expect personalized, integrated experiences. Regulators expect real-time transparency. Shareholders expect faster growth and greater operational productivity. Without a fundamental shift in how data is managed, financial institutions will be unable to meet these heightened expectations and unlock the full potential of AI.

As the industry moves deeper into this new era, banks that manage data effectively will gain a significant innovation and growth advantage. Many institutions still struggle with siloed data and how to integrate disparate data types – structured, semi-structured, and unstructured – into their data architecture. Even when data is centralized, making it efficiently available and reusable across a large, complex organization operating in a highly regulated environment presents a persistent challenge.

In this paper, we examine how building a connected data foundation with accessible, reusable data products can mitigate these obstacles. We identify the critical factors that leaders should consider when assessing next-generation data platforms and implementation partners, and we analyze the experiences of two financial services organizations putting such a platform into practice.

As these case studies demonstrate, the data opportunity is significant. An effective model, implemented well, can deliver smarter, more personalized intelligence that accelerates innovation and cuts time-to-market for new or tailored offerings. It can solve specific problems across domains as broad as compliance reporting, sales, and credit decisioning, unlocking AI and wider digital transformation and accelerating data-driven advantage.



Banking's Data Imperative

Deploying effective data solutions is becoming non-negotiable as evolving customer needs and accelerating agentic and generative AI capabilities rapidly reshape the financial landscape. The ability to manage data well is an increasingly key differentiator for institutions aiming to compete and grow in this fast-paced industry.

However, while digital advancements have led to impressive efficiency gains across the sector, based on Cognizant's findings, the perceived gains from digital transformation may be less than they seem. Automation, for instance, dramatically cuts transaction processing costs by 60–70% but also introduces a 12–15% rise in back-office operational costs from data management overhead. And while digital infrastructure has fuelled 10–50x growth in some markets, it's also constrained innovation by requiring significant custom API and middleware investments, which consume 15–20% of IT budgets.

Similar data integrity issues are prevalent across the organization. Inconsistent data has triggered a 15-20% increase in regulatory inquiries due to audit reconciliation errors. Incompatible risk scoring models have inflated the time spent compiling risk assessment reports by 40%, while siloed data has resulted in a 30-35% expenditure on ineffective marketing campaigns.

Supporting the business growth message:

3-5x

increase in back office productivity

10-50x

segment growth

3-5x

increase in AUM per advisor

Supporting the data issues message:

12-15%

rise in back office data management overhead

15-20%

IT budgets on custom APIs and middleware

15-20%

more annual audit error-related regulatory inquiries

Two Data-Driven Barriers to Innovation and Growth

To accelerate innovation and growth, banks must solve the twin digital challenges that lie behind many of these barriers to digital productivity: data fragmentation and data accessibility.

Data Fragmentation: Despite significant investments in enterprise data strategies, many banks are still burdened by unconnected data spread across disparate legacy systems and in forms that are costly to process or automate. High volumes of unstructured data – such as earnings transcripts, customer agreements, and regulatory compliance-related records in PDF format – are particularly challenging for ingestion and analysis. This data fragmentation acts as a brake on innovation. It prevents banks from cost-efficiently sharing intelligence, thereby making it difficult to deploy advanced analytics and AI innovation at scale.

Data Accessibility: Even with robust data storage and data integration, leveraging centralized data for multiple use cases across a highly regulated global financial institution requires significant investment in data engineering, security, and governance. To close the digital productivity gap and drive faster innovation, data accessibility needs to be easy to manage and integrate into solutions that solve specific business problems.



The Path to Value: From a Unified Foundation to Accelerated Delivery

Overcoming the barriers of unified and accessible data to unlock growth and innovation is a three-step process. First, build a unified data foundation. Then, use it to create reusable data products that are easily accessible organization-wide. Finally, accelerate time-to-value and mitigate the risks of complex transformation with effective solution design and delivery - where trust and governance are integral to how the solution is architected and executed, and an important part of the process.

When delivered effectively, such a solution can transform raw data from client interactions across the bank into a high-value asset for strategic decision-making. It can turn a single source of truth on customer data into varied outcomes—from reducing marketing costs to preventing fraud more effectively and improving client risk management.

The Foundation: Unified, Scalable Data

Cloud and AI technologies are transforming how banks connect data across disparate locations and business units. They also unlock the value of unstructured data, which traditionally was too expensive and error-prone to process manually.

A good example of this is a SaaS platform for integrated electronic front-office trading, portfolio management, and risk management which spent 4,000 hours processing an estimated 100,000 emails—many of which were critical notifications that could cause production outages if missed. By automating this intake with the Snowflake AI Data Cloud and Snowflake's Cortex AI capabilities, the company eliminated all missed notifications.

Beyond basic ingestion, an effective data repository must highly automate ETL processes to reduce management overhead. Streaming pipelines, for example, are critical for real-time needs such as fraud prevention, market analysis, or digital product approvals. To meet regulatory and resilience standards, the architecture must also be structured, scalable, and secure, supporting multi-cloud environments.

The Journey to Unified, Scalable Data

Transitioning to a unified data foundation requires careful change management. A phased approach ensures the end state is robust and ready to deliver the desired business value.

Category	Initiate (Setup Phase)	Experiment (Pilot Phase)	Evolve (Future State)
Data Architecture	Adopt initial architecture to minimize disruption to downstream systems.	Pilot evolved architecture to compare efficiencies by removing redundant data layers.	Full pivot to new architecture across teams to maximize efficiency.
Capabilities & Frameworks	Catalog basic data quality and establish access management.	Test enhanced frameworks through pilot programs to improve data quality and governance.	Consolidate best capabilities and implement across future initiatives.
Team Operating Model	Assign domain ownership to teams to initiate progress.	Identify need for specialized domain expertise for expanding use cases.	Transition to expertise-based teams for better scalability and focus.
Vendor Model	Accelerate development and test collaboration models using a multi-vendor ecosystem.	Compare pilot programs for single vendor vs. multi-vendor efficiencies.	Shift to end-to-end vendor model for streamlined delivery and efficiency.
Measurable Outcomes	Focus on execution rather than measurement.	Introduce basic KPIs to monitor efficiency and lessons learned.	Establish robust processes to define and track KPIs like turnaround time and usage.
Platform	Introduce a modern tech stack to replace legacy systems.	Add/remove tools based on pilot learnings (e.g. data virtualization, BI tools).	Finalize optimized tech stack aligned with business needs and operational learnings.

The Value Creator: Accessible, Reusable Data Products

To truly unlock the value of unified data, banks need to go beyond centralization and focus on creating reusable data products. These products should be easy to find, understand, and apply to real business problems.

Simplifying the data engineering process accelerates time-to-value by empowering a broader range of data analysts and scientists to participate. This is where composable, modular elements are effective. Platforms such as Snowflake are designed to deliver this capability out-of-the-box, also enabling developers to develop more complex data pipelines or machine learning models directly on the platform. When developers can work with the data where it's stored, they save valuable time, minimizing the risks of moving data and broadening the scope for innovation.

To turn these capabilities into broad business value, discoverability must be just as simple. For example, if institutional data and regulatory requirements are curated into data products that are difficult for teams to find, evaluate, or integrate, they won't deliver the desired efficiency or accuracy improvements.

Integrating a data marketplace into your data solution makes the discovery and evaluation process easy. Users of the Snowflake Marketplace, for instance, benefit from robust search and filtering capabilities, along with detailed descriptions and sample data. This not only reduces demands on central IT teams but allows the business to innovate faster.

What is a Data Product?

A mature data product isn't just a dataset—it's a well-packaged asset with qualities that make it usable across the organization. Think of it like a well-labeled, secure, and trusted ingredient in a recipe:

- Findable in a catalog.
- Accessible through defined APIs.
- Interoperable with other systems.
- Reusable without needing extra cleanup.
- Secure and governed.
- Trustworthy with clear lineage.
- Owned by someone accountable.
- Self-describing with embedded metadata.
- Valuable to its consumers.

Embedding these traits into your data products ensures they're not just technically sound, but also business-ready—empowering teams to innovate faster and with confidence.



The Value Accelerator: Solution Design and Delivery

For global financial institutions undergoing constant transformation, and balancing AI innovation with high expectations, minimizing delivery fatigue and resource-intensive mistakes are top priorities. Leaders seeking an integrated data solution must choose partners with deep, domain-specific solution design and delivery expertise that mitigates these risks and ensures success. Critical capabilities to look for include:

- **Demonstrated Results:** Proven experience in designing and delivering global data solutions maximizes the likelihood of meeting strict security, governance, cost, and business outcome requirements.
- **Delivery Accelerators:** A deep bench of certified experts, coupled with industry-specific solutions, and global offerings will accelerate delivery within financial services' regulatory and operational constraints.
- **Business Value Focus:** A holistic approach to data to drive business value and enable AI at scale that extends beyond technical implementation ensures the focus remains on how data can be leveraged to drive business value and enable AI at scale.

Value Hotspots: Driving Results with Secure, Trustworthy, Accessible Data Products

- **Fraud Detection:** Real-time anomaly tracking enables immediate fraud detection and prevention. It also powers the automation of Anti-Money Laundering (AML) and Know Your Customer (KYC) processes, reducing manual effort and speeding up compliance checks.
- **Regulatory Compliance:** The ability to access a single source of truth streamlines and automates reporting for regulatory bodies such as the Financial Conduct Authority (FCA), while also optimizing audits through easily accessible and well-governed data.
- **Customer Insights:** AI-driven personalization and financial recommendations allow banks to offer more relevant products and services that improve the customer experience.
- **Credit Decisioning:** With a complete view of customer data, banks can build more accurate AI-based lending models for faster, more confident credit decisions and more precise risk analysis.



Putting it into Practice: Enhanced Customer Intelligence

A leading financial services organization faced the challenge of a legacy data infrastructure that hindered its ability to effectively serve a large customer base of small business and commercial clients. They needed a next-generation data platform to reduce costs, accelerate time to market, and deliver a real-time, 'always on' online and digital experience.

To solve this, the organization partnered with **Cognizant and Snowflake** to build a single, unified enterprise data hub on the Snowflake AI Data Cloud. This collaborative effort established a robust and scalable digital foundation that created a single source of truth for near real-time customer data. By leveraging state-of-the-art DevOps automation, the project achieved a fast and accurate software delivery process, leading to a 40% annual saving in licensing and infrastructure costs.

This foundation of real-time customer intelligence, built with Cognizant and the Snowflake AI Data Cloud and accelerators, unlocked significant business value across the organization, including:

Faster, More Cost-Effective Credit Decisioning:

- New credit models are now built in 6-8 months instead of 2-3 years.
- Decades of credit data can be aggregated in hours, not weeks.
- Processing time was reduced by 80%, leading to a \$57M annual saving.
- Credit models on the cloud are now 2x faster.

More Effective Digital Marketing:

- New social media customer acquisitions grew from 2% to 6%.
- Cross-channel marketing spend was cut by 26%.

Improved Fraud Prevention:

- Voice biometrics were enabled using Snowflake machine learning capabilities.
- Transaction decisioning was improved with Snowflake Direct Data Share.



What Drives Success? Critical Factors for Next-Generation Data Platforms

While every bank faces unique opportunities and challenges when modernizing its data foundation, successful projects consistently share the same core features. As business and technology leaders evaluate potential solutions and partners, they should focus on these critical success factors, that ensure the foundation is scalable, governed, and ready to support next-generation capabilities like AI

Strategic Partnerships and Value

Domain Expertise and Delivery Experience: Working with a partner who combines deep, domain-specific experience with certified technical expertise can significantly de-risk your project. This ensures accelerated delivery and a faster return on investment.

Maximizing Value Realization: A centralized data platform should be more than just a data repository; it should be a catalyst for new opportunities. The most valuable platforms include built-in capabilities including AI-driven insights, access to external data marketplaces, and the ability to rapidly develop new data products.

Security and Governance

Secure by Design: A solution must be built with industry-leading, cybersecurity-first principles at its core. Partnering with a provider who embeds these principles ensures your sensitive data is fully protected from ever-evolving cyber threats.

Centralized Data Governance and Control: To de-risk AI from a regulatory and ethical standpoint, all AI features—including agentic AI—must operate within the same centralized governance plane. This ensures consistent controls across your entire data landscape, from raw data to advanced AI insights.

Technical Foundation and Performance

Comprehensive Data Ingestion: A robust data foundation must be capable of efficient data ingestion across unstructured, semi-structured, structured forms. The right solution enhances data quality and accuracy while also optimizing cost.

Scalable Real-Time Analytics: To handle the demands of modern banking, platforms must support low-latency data ingestion and have a multi-cluster architecture. These features ensure smooth processing and fast data updates, even during high-volume periods or workload spikes.

Cost-Efficient Data Management: The most effective solutions use a high degree of automation and a scalable architecture to optimize compute, storage, and data loading. Look for robust resource monitoring capabilities that give your teams the control they need to manage and reduce costs.



Putting it into Practice: Strategic Client Decisioning

A leading bank sought to optimize the value of its client relationships by calculating a series of key scores and making them available to teams across the bank as a data product. This initiative required integrating data from every touchpoint across a complex financial architecture to create a unified view of each client.

To address this challenge, **Cognizant and Snowflake** partnered with the bank to develop a **medallion architecture**. This framework was designed to ingest and process data from diverse source systems in a structured, scalable, and secure manner. The architecture consists of three distinct layers:

- **Bronze Layer:** This ingestion layer collects raw, unprocessed data from all source systems, serving as a complete and unaltered repository.
- **Silver Layer:** In this processing layer, data is cleaned, standardized, and conformed to a consistent format, ensuring high quality and accuracy.
- **Gold Layer:** This final, curated layer is where the client relationship value scores are calculated. The data here is aggregated and made business-ready, serving as the single source for reporting and analytics.

The three-layer model progressively refined the data while preserving its quality and lineage at every stage. The resulting client relationship scores weren't just calculated—they were published as data products designed to be discoverable, secure, and self-describing. This meant that teams across the bank could easily access and trust the data, knowing it was governed, high-quality, and aligned with business goals. The data products could also be consumed directly by business intelligence tools and downstream applications.

It's this kind of thoughtful design that turns data into a strategic advantage. Access to the client relationship scores has enhanced strategic decision-making and unlocked new value including:

- **Optimizing Product Sales:** Teams can now identify cross-sell and upsell opportunities with precision.
- **Proactive Client Risk Management:** The bank can monitor client relationships in real time, allowing for more proactive risk mitigation.
- **Value-Based Team Performance Measurement:** Relationship managers can now be evaluated on the value they create, aligning incentives with strategic goals.



A Model for Data-Driven Advantage

Digital banking has already established data as a strategic asset, a trend now being rapidly accelerated by AI-driven opportunities. A connected and trusted data foundation is crucial for any bank with agentic AI ambitions: AI agents cannot act intelligently or make relevant decisions without it.

Beyond the AI opportunity, unified data is critical for wider digital transformation. Banks that integrate such a foundation with reusable data products that are easily accessible across their organization will be positioned for significant business benefit—from faster innovation and smarter decision-making to greater operational efficiency.

Executing such large-scale data transformation across a global bank is a complex undertaking. Leaders must carefully consider whether they have the right domain, solution, and delivery expertise to de-risk the project and maintain a sharp focus on driving business value.

The **Cognizant and Snowflake** partnership provides a powerful solution to this challenge. It combines Snowflake's next-generation AI data platform with the deep experience, accelerators, and frameworks that Cognizant brings as a proven data transformation and delivery partner. As long-standing strategic partners, their close working relationship is built on a shared focus on AI and data transformation that accelerates a decisive data-driven advantage.

Key strengths of their AI-powered solution include:

- **Cortex Analyst:** Enhancing structured data analysis and layering on top of the Snowflake Data Platform to give the end user AI-powered SQL generation ability with semantic model understanding.
- **Cortex Search:** Natively providing a high quality search engine for all varieties of data, combining vector and lexical (keyword) search with an additional semantic reranking step to deliver high-quality, low-latency retrieval at scale.
- **Snowflake Intelligence (agentic orchestration layer):** Connecting to and querying data from multiple sources across the enterprise, including data lakes, databases, data warehouses, and even other lakehouse catalogs.
- **Snowflake Marketplace:** Centralizing data access so AI agents can easily work with whatever data the user wants, from historical data for training and pattern recognition to operational data for business critical decision-making across lines of business.

Unlock Your Data-Driven Advantage with Cognizant and Snowflake

- Award-Winning Partnership
- Deep Industry Expertise
- AI-Driven Solutions
- End-to-End Implementation
- Demonstrated Results

Discover how to unlock your data advantage with Cognizant and Snowflake.



© Copyright 2025–2027, Cognizant. All rights reserved. No part of this document may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the express written permission of Cognizant. The information contained herein is subject to change without notice. All other trademarks mentioned herein are the property of their respective owners.