

Real-Time Intelligent Applications: The Art of the Possible

Data Platform Requirements for Intelligent Operational Applications

WHITE
PAPER

 | VENTANA RESEARCH

Sponsored by:

DATASTAX



Table of Contents

| | |
|---|----------|
| Understanding The Need for Data-Driven Business Operations | 3 |
| Unlocking the Value of Data | 3 |
| Identifying Data-Driven Technology Requirements | 5 |
| Addressing Common Challenges | 6 |
| Making the Business Case for Investment..... | 7 |
| About Ventana Research | 9 |



Understanding The Need for Data-Driven Business Operations

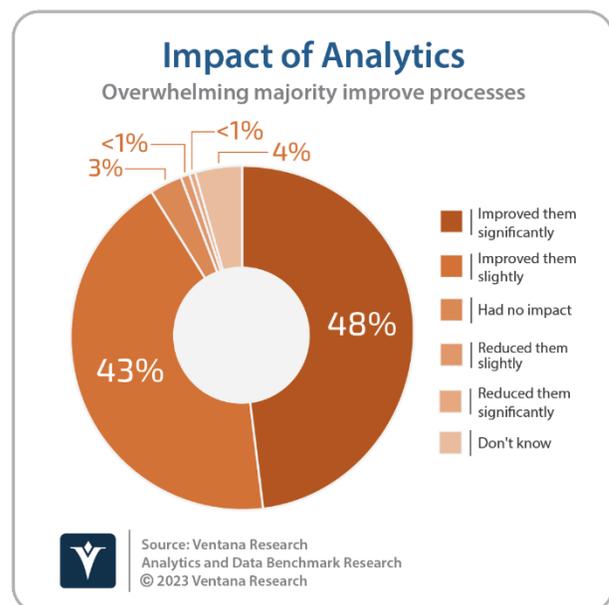
In the current economic climate where customers expect more from the organizations with which they do business, companies are looking to be more responsive and improve efficiency while cutting costs. Data is a fundamental enabler of new business processes and applications, allowing organizations to deliver contextually relevant recommendations, predictions and forecasting to customers and workers, increasing efficiency and improving customer service.

All businesses utilize at least some form of data processing and analytics, but the most successful organizations build an entire culture around it. Data-driven organizations stand to gain competitive advantage, responding faster to worker and customer demands for more innovative, data-rich applications and personalized experiences. Although many organizations may aspire to be more data-driven, identifying and defining the steps required to achieve that goal are not necessarily easy.

The shift to more agile business processes requires responsive data platforms and applications that are driven by machine learning. Consumers are increasingly engaged with data-driven services that are differentiated by personalization and contextually relevant recommendations. Additionally, worker-facing applications are following suit, targeting users based on their roles and responsibilities. The pandemic—quickly followed by a challenging economic climate—has emphasized that data-driven organizations stand to gain competitive advantage, responding faster to worker and customer demands for more innovative, data-rich applications and personalized experiences.

Unlocking the Value of Data

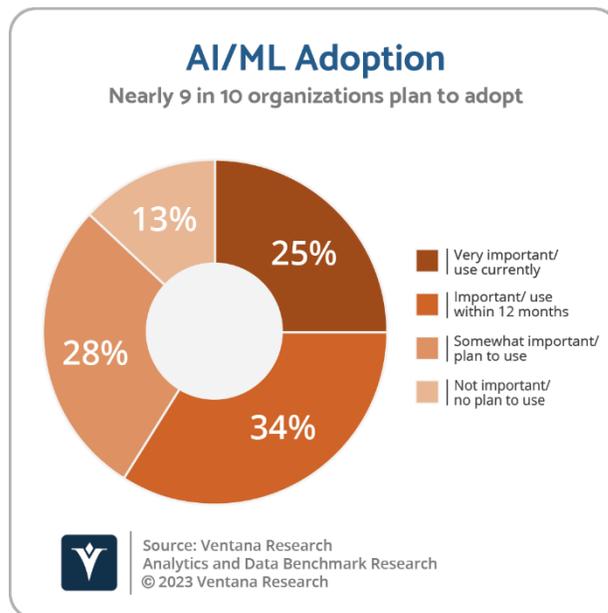
Organizations have long understood the value of data. An overwhelming majority of businesses (91%) indicate the use of analytics has improved operational processes, with almost one-half (48%) seeing significant improvement. Business benefits cited in our research from analytics are improved communication, increased competitive advantage, improved customer experience and satisfaction and improved management and alignment of business.





When analytics tools incorporate artificial intelligence and machine learning capabilities, they can be more readily used by line-of-business professionals throughout the organization. It is no surprise, then, that nearly 9 in 10 organizations either use or plan to adopt AI/ML technologies, with more than one-half citing usage now or plans to adopt within a year.

Consumer-facing interactive applications driven by machine learning recommendations can be found in nearly every corner of the market. For example, consider the needs of an online eyewear vendor. As with most online retail applications, customers expect a personalized shopping experience with real-time product recommendations and, in some instances, the ability to try on products virtually. Data from past purchases and current browsing behaviors must be combined to ascertain a customer's preferences, budget and items likely to be of interest.



Organizations learned during the pandemic that they needed to be more responsive and adaptable to evolving business conditions and customer demand. Retailers needed to adapt quickly to an increase in online ordering, home deliveries and curbside pickup. Those that pivoted their consumer-facing applications and business processes were able to meet demand for digital platform sales. Maintaining customer service levels requires flexible supply chains to move and deliver products from retail stores and warehouses, ensuring efficient use of stock levels. This flexibility is enabled by intelligent applications that provide decision-makers with accurate and timely data related to stock levels, logistics and fulfillment capabilities as well as predictions and forecasting about customer demand.

Increasingly, workers expect similar levels of personalized and contextually relevant content as well as responsiveness in employee-facing applications. Human capital management is a prime example. HCM leaders have typically used dashboards filled with visualizations that provide snapshots of historical metrics with little or no context as to what each data point means to strategic objectives and initiatives. Our research shows that an increasing number of organizations are using or planning to deploy people analytics solutions to help predict and manage workforce-related challenges and opportunities. In fact, we assert that by 2024, one-half of organizations evaluating people analytics offerings will require prescriptive guidance with actionable insights for quantifiable HCM impact.

Where data scientists were once the only group required to synthesize and interpret data for business impact, that power now increasingly sits directly with the individuals responsible for taking action. Increased use of AI in the office of finance has enhanced the breadth of analytics used by financial planning and analysis organizations, applying machine-driven task



supervision to speed process execution and reduce errors, generating recommendations and automating report commentary. Ventana Research asserts that by 2025, almost all vendors of software designed for finance organizations will incorporate some AI capabilities to reduce workloads and improve performance.

Increasingly, software vendors targeting the office of finance will differentiate offerings by the capabilities and accuracy of AI functionality. And organizations will adopt this technology to attract and retain the best talent, because AI enables substantial reduction of low-value work. This means there is more time for workers to focus on tasks that require their expertise, experience and judgment.

Identifying Data-Driven Technology Requirements

We know data-driven organizations when we see them—Airbnb, DoorDash, ING Bank, Netflix, Spotify and Uber are often cited as examples. But it is not always clear to those aspiring to that level of success what separates data-driven organizations from the rest. While these organizations often have many workers and large budgets, the key difference is the ability to personalize and improve the customer experience through the delivery of contextually relevant recommendations in real-time.

Our Analytics and Data Benchmark Research shows that only 18% of organizations can be considered part of the innovative top tier when it comes to analytics and data usage. Organizations of any size can and should apply the same level of thinking and execution to data-driven applications and decision-making, and that starts with adopting real-time analytics.

The need for real-time interactivity has significant implications for data platform functionality required to support these applications. To protect the performance of the operational workload, traditional architectures have involved the extraction, transformation and loading of data from the operational data platform into an external analytic data platform, enabling the operational and analytic workloads to run concurrently without adversely impacting each other.

Operational Data Platforms

Market Assertion

Through 2026, operational data platform providers will continue to invest in hybrid operational and analytic processing capabilities to support growing demand for data-intensive intelligent operational applications.



VENTANA
RESEARCH

Matt Aslett
VP & Research Director

AI-driven intelligent applications require a new approach that enables real-time performance of machine learning on operational data to deliver instant, relevant information to accelerate



business decision-making. We assert that through 2026, operational data platform providers will continue to invest in hybrid operational and analytic processing capabilities to support growing demand for intelligent operational applications infused with personalization and AI-driven recommendations.

Intelligent applications, while operational in nature, rely on real-time analytic processing to deliver functionality, including contextually relevant recommendations, predictions and forecasting. Real-time analysis of data is still relatively rare in most organizations, with less than one-quarter (22%) of participants in Ventana Research's Analytics and Data Benchmark Research reporting that their organization analyzes data in real time.

To generate the maximum value from real-time data, organizations need to process events as they occur, but also analyze these events in the context of historical data to forecast the likely impact and recommend potential responses. Many organizations have invested in event



Data-driven organizations have holistic environments that combine real-time and historical data, as well as machine learning.

processing and streaming analytics alongside traditional batch data processing, with 30% citing usage of streaming data for more than one year. However, in a lot of cases these have been treated as two distinct approaches that are not well integrated. In comparison, data-driven organizations have holistic environments that combine real-time and historical data, as well as machine learning. While ML model training may well still be performed on historical event data streamed to an external analytic database, storing the resulting features and performing model serving and inference on real-time event data in the operational database supports the delivery of intelligent operational applications.

Data-driven organizations have also typically been early adopters of cloud-managed services. Established organizations are adopting cloud services at a rapid rate while balancing existing investments in on-premises infrastructure. As such, there is increased interest in data platforms that span on-premises and multiple public clouds.

Addressing Common Challenges

The state of being data-driven is achieved through a combination of people, processes, information and technology improvement. But high costs, complexity and scaling issues have been roadblocks for many organizations in achieving dynamic, real-time intelligence in operational platforms. Data, especially in large volumes, is complex and inherently costly. Our latest research reveals that only 59% of organizations feel very confident or confident in the ability to analyze accumulated data. Additionally, the skills necessary to effectively use data are often lacking, requiring investments in training to ensure that people at all levels of an organization can understand and work with the data.



Training machine learning models to produce accurate and highly relevant results requires massive volumes of historical data. Moving large volumes of data between duplicate platforms is costly and complex, and yet existing approaches rely on batch movement of operational data into an external data warehouse or data science platform to train and



Leaders are increasingly bringing machine learning inferencing to the data, enabling real-time online predictions and recommendations.

execute a model before pushing the results back into the production application. While data-driven companies continue to use specialist analytic and data science platforms to train models offline, leaders are increasingly bringing machine learning inferencing to the data (rather than taking the data to the ML models), enabling real-time online predictions and recommendations.

Scalability of data platforms has traditionally focused on scaling up to deliver additional processing power in response to growing data and query volumes. This vertical scaling can be expensive and does not support AI-infused applications with distributed low-latency performance requirements. Scaling out horizontally, however, supports business continuity and enables organizations with globally distributed customers and

workers to deliver applications that are available and performant at all times, in all geographic regions. This has driven the adoption of scale-out distributed SQL and NoSQL databases by many data-driven organizations. We assert that through 2025, two-thirds of organizations will re-examine their current operational database suppliers with a view to improving fault tolerance and supporting the development of new intelligent operational applications with scale-out databases like Apache Cassandra, Redis and others.

Managing and governing data across multiple regions and data centers is a key consideration, especially for remaining compliant with data security, data privacy and data sovereignty requirements. Governance is increasingly important in relation to the use of machine learning models, particularly the ability of organizations to ensure that models are ethical, explainable, trusted and compliant with organizational policies or regulatory requirements. Time-based model creation facilitates good governance. As organizations continually run models, the ability to attribute time to model inferencing is increasingly important for organizations to travel back in time to understand and demonstrate for regulatory purposes how automated predictions, recommendations and decisions were made at any given point in time.

Making the Business Case for Investment

Organizations that ignore the call for real-time data-driven decision-making risk falling behind in the market. Some will be more successful than others. Those that are successful approach overall data architecture as a single environment rather than operational and analytic silos.



Operating in real-time environments is also key to enabling organizations to be more flexible, agile and responsive and eventually improve customer satisfaction and efficiency.

A combination of people, processes, information and technology improvement are required to enable real-time, data-driven business operations. There are countless data and analytics products and services available that promise to accelerate and improve data processing and analysis. However, the best products in the world will not help if an organization does not also take steps to improve and refine its culture and business processes as well as the information collected and used to inform decision-making. Almost 1 in 5 participants (19%) in Ventana Research's Analytics and Data Benchmark Research said a lack of executive support is a barrier to making improvements to analytics and data. The support of executives and business leaders is critical in defining, articulating and demonstrating the values, vision and goals that promote a culture of data-driven optimization and decision-making.

We recommend that organizations think about data architecture holistically. Most organizations have separate operational data platforms, ML platforms and streaming platforms. Ideally, there would be a unified solution that delivers all of these. Even if a single data stack does not yet exist, reducing complexity should be a goal of the

organization. Our research leads us to assert that, through 2025, data platform and data management vendors will continue to prioritize support for data streaming and event processing functionality, adapting products to address the growing demand for real-time data processing. Organizations should evaluate potential suppliers on the ability to deliver on a combination of those capabilities, whether delivered as one product or several products.

Intelligent applications that provide contextually relevant recommendations, predictions and forecasting to customers and employees will soon be the norm. To be successful, organizations need to adopt data platforms and business processes and ensure support for the expectations of customers and employees alike in terms of real-time processing and interactivity.

Streaming Data and Events
Market Assertion

Through 2025, data platform and data management vendors will continue to prioritize support for data streaming and event processing functionality as they adapt their products to address the growing demand for real-time data processing.

Matt Aslett
VP & Research Director



About Ventana Research

Ventana Research is the most authoritative and respected benchmark business technology research and advisory services firm. We provide insight and expert guidance on mainstream and disruptive technologies through a unique set of research-based offerings including benchmark research and technology evaluation assessments, education workshops and our research and advisory services, Ventana On-Demand. Our unparalleled understanding of the role of technology in optimizing business processes and performance and our best practices guidance are rooted in our rigorous research-based benchmarking of people, processes, information and technology across business and IT functions in every industry. This benchmark research plus our market coverage and in-depth knowledge of hundreds of technology providers means we can deliver education and expertise to our clients to increase the value they derive from technology investments while reducing time, cost and risk.

Ventana Research provides the most comprehensive analyst and research coverage in the industry; business and IT professionals worldwide are members of our community and benefit from Ventana Research's insights, as do highly regarded media and association partners around the globe. Our views and analyses are distributed daily through blogs and social media channels including [Twitter](#), [Facebook](#) and [LinkedIn](#).

To learn how Ventana Research advances the maturity of organizations' use of information and technology through benchmark research, education and advisory services, visit www.ventanaresearch.com.