CLOUDERA

EBOOK

The Retail and eCommerce AI Playbook

Winning customers in the age of digital commerce with real-time personalization powered by generative and agentic AI.

Table of Contents

Chapter 1:
The Need for Personalization in Retail

Chapter 1:
The Need for Personalization in Retail

Retailers are under market, competitive, and regulatory pressure to transform. As the line between digital and physical shopping experiences continues to blur, and as customers have more options available and more companies targeting them, it is becoming increasingly difficult to acquire customers and increase loyalty and share of wallet.

Customer Expectations: The Amazon Effect

Customer expectations are often set by best-in-class digital platforms built by big technology firms, not by direct competitors. Customers have been conditioned to expect proactive, personalized service that feels frictionless.

Relevance Is Critical. Customers are increasingly willing to share their data, but in exchange, they expect a tailored and relevant experience. According to McKinsey & Company, 71% of consumers expect companies to deliver personalized interactions, and 76% get frustrated when that expectation is not met.

Competitive Pressures: Rising Above the Noise

The retail space has never been more crowded. With the rise of direct-to-consumer brands and the dominance of global marketplaces, traditional retailers can no longer compete on price or inventory alone.

Experience Matters. With endless purchasing options just a click away, experience loyalty is replacing brand loyalty. According to a Qualtrics survey, 51% of consumers reduce or stop spending with a company after a poor experience.

Efficiency and Margins. Advertising costs are increasing as more brands are vying for digital shelf space, increasing bidding and raising the cost per click. It's more important than ever to target the right customer with the right offer at the right time to optimize acquisition costs.

Omnichannel Retail Strategies: Closing the Loop

The modern customer journey is no longer linear. It spans multiple channels and platforms, including social media, websites and applications, and physical stores. The challenge for retailers is maintaining a single thread of conversation with each customer throughout their journey, regardless of where it occurs.

Omnichannel Engagement Increases Revenue. Marketing to customers through multiple channels can increase purchase rates by 287%.

AI has the potential to support personalization and build a single thread of conversation across channels, but only if it is built on a unified, consistent view of customer data.

Chapter 2: The Personalization Challenge

Chapter 2:
The Personalization Challenge

Retailers have been on the journey towards personalized omnichannel engagement for a decade or more, and yet they are still struggling. There are often multiple technical and organizational barriers between knowing the customer and taking proactive action based on that knowledge.

Siloed and Distributed Data. Most retailers are managing a complex collection of siloed systems acquired through years of software acquisitions, expansions, and replatforming efforts. These systems often lack interoperability, so customer data is fragmented, and it's difficult to find, access, integrate, and analyze. A unified, consistent view of the customer that can be delivered and leveraged across departments and channels is critical for personalization, and it's also a massive undertaking.

Data Volume and Variety. In the past, customer data was often limited to transaction history. Today, customer data can include clickstream logs, social media sentiment, customer service call transcripts, and even video analytics from physical stores. It includes structured and unstructured data at a scale that can often overwhelm traditional data management systems.

The Need for Real Time. Many retailers still rely on batch processing to ingest data, limiting their ability to capture the moment when the customer is most likely to make a purchase. By the time these retailers have successfully processed and identified a customer's intent signals, they have likely already moved on. Moving from historical to real-time context requires real-time ingestion, processing, and automated decision making at the edge.

The AI Talent Gap. While customer expectations for personalization are driven primarily by the biggest technology firms in the world, retailers often lack the data science talent required to operationalize AI and deliver real-time personalization at scale.

To overcome these challenges, retailers need to rethink their data architecture. They need to build a foundation of trusted customer data, build AI models that can leverage that data to personalize the customer experience across channels, and deliver those experiences in real time to drive more revenue.

Chapter 3:
Building the Data Foundation for a Customer 360

Chapter 3:
Building the Data Foundation for a Customer 360

A Customer 360 is crucial for personalization in retail. In many cases, a failed attempt at personalization is worse than no attempt at all: without a unified, consistent view of the customer across channels and touchpoints, retailers not only lose out on revenue; they risk alienating customers altogether.

Building a Customer 360 starts with building a foundation of trusted data. It requires bringing distributed data together, ensuring consistency and quality, and making it accessible for analytics and AI.

Unified Data. Unified data is the first step in building a Customer 360. This does not mean migrating data to a single system, cloud, or data center, which would be virtually impossible for most retailers. Instead, it means implementing a data fabric that can provide governance, security, discovery, and access to a consistent view of the data across environments.

Consistent Data. Consistency is the key to maintaining a single thread of conversation across channels and touchpoints, and supporting a nonlinear customer journey. Multiple copies of data across departments and channels make it difficult to understand where a customer is in their journey, and incorrect assumptions can lead to customer friction and lost opportunities.

Open Data. One reason building a Customer 360 has been historically difficult in retail is the proliferation of point solutions and the adoption of proprietary data platforms and tools for working with data. With the rapid expansion of AI tools, retailers cannot afford for their data to be locked into a proprietary format or system. Customer data must be easily accessible via open table formats like Apache Iceberg, and maintained within a consistent security and governance framework in order to be leveraged for AI.

To solve these data challenges, retailers should modernize their data architecture by adopting an open data lakehouse that combines the flexibility and scalability of data lake storage with the governance, security, and data management capabilities of a data warehouse. An open data lakehouse provides a consistent framework for collecting, processing, securing, and accessing customer data across distributed environments, while also giving data teams and consumers the freedom to choose from a variety of tools and engines for working with the data for analytics and AI use cases.

Chapter 4: AI-Powered Insights—Agentic Workflow Built with Cloudera

Chapter 4:
AI-Powered Insights—Agentic Workflow Built with Cloudera

A Customer 360 is the foundation for real-time personalization, but the critical next step is to move from descriptive to predictive and prescriptive analytics. AI is the technology that transforms customer insights into differentiated, revenue-generating experiences by understanding the intent and context behind each interaction and taking the appropriate action. For retailers, this means moving beyond machine learning models towards generative and agentic AI.

Generative AI Provides Relevance at Scale. Generative AI (GenAI) gives retailers the ability to quickly understand customer context by automatically synthesizing massive volumes of structured and unstructured data from across departments and channels, and making intelligent assumptions about where that customer might be in their buying journey. That context can empower AI chatbots, store employees, and customer service agents to confidently deliver the next best action, product, or conversation.

Delivering AI Insights Using Natural Language. For retailers, a majority of the employees who need access to data for customer insights are non-technical. They cannot write SQL queries or build models themselves. GenAI closes the talent gap, democratizing insights by providing a natural language interface for employees and even customers to be conversational with their data. This critical self-service capability ensures a consistent view of the customer and a single thread of conversation across the customer journey.

Agentic Workflow Built with Cloudera

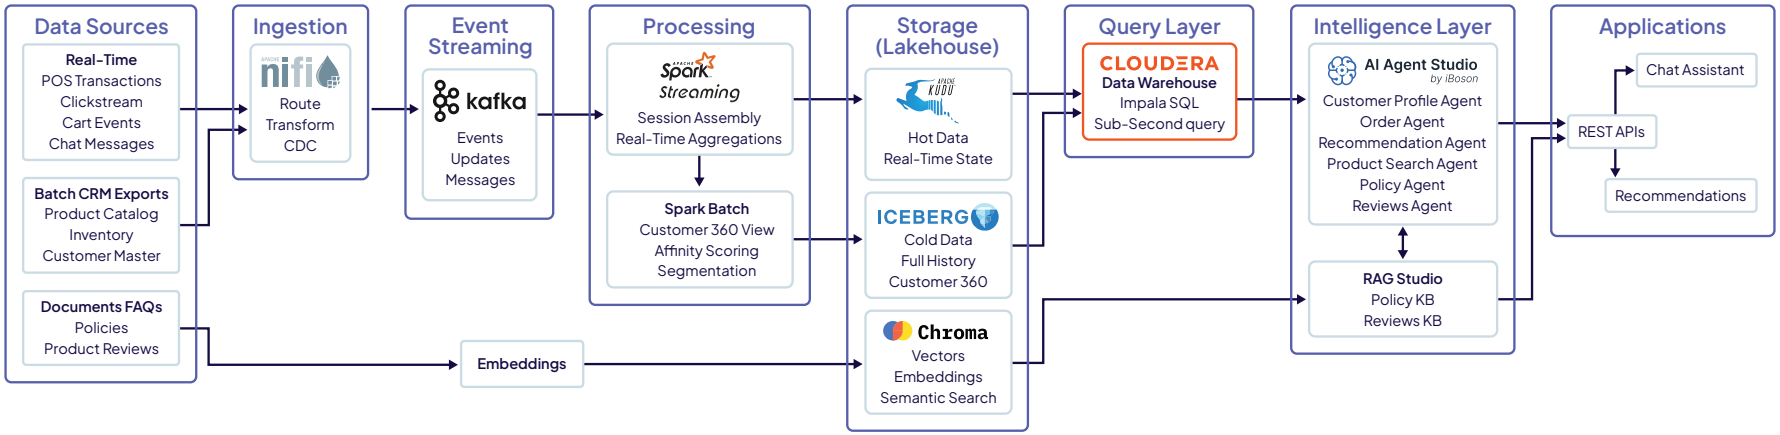

Agentic AI Turns Insights into Action. Agentic AI enables the automation of customer interactions, taking customer insights and executing tasks on behalf of the customer and the business. From a digital concierge who can act as a personal shopper for customers to an agentic customer service agent who can investigate complaints issues and take remedial action, agentic AI has the potential to truly transform the customer experience.

Multi-Agent Orchestration unlocks the full potential of AI-powered personalization. In a multi-agent system, a manager agent breaks down a goal into smaller sub-tasks and assigns them to specialized agents to complete. This strategy solves three big challenges for retailers pursuing AI-powered personalization.

- **Contextual Continuity:** By sharing a unified Customer 360 data layer, every agent understands the customer's entire history, preventing a common issue in traditional support where the customer feels like they are starting over with each new representative.
- **Autonomous Problem Solving:** Multi-agent systems take a goal and break it into sub-tasks, working across departments without requiring human intervention.
- **Scalable Personalization:** Retailers can provide a personalized concierge for every customer, scaling to thousands of unique, concurrent workflows tailored to individual shopping behaviors.

Putting It All Together: An Autonomous Exchange Workflow

To see how these capabilities function together, consider an AI-powered Customer 360 workflow for handling a customer request for an exchange. Instead of a single, rules-based chatbot, we will use a multi-agent orchestrator and a team of specialized agents, each querying live enterprise data to address the request:

01 **The Profile Agent** instantly pulls the customer's profile using the data fabric to verify loyalty status and lifetime value.

02 **The Policy Agent** reasons through the retailer's current return and compensation guidelines to determine what options are available.

03 **The Recommendation Agent** cross-references the customer's style preferences with live inventory to suggest an alternative item that is currently in stock.

04 **The Fulfillment Agent** places the order once the customer confirms and provides order and tracking information.

By leveraging multi-agent orchestration and a Customer 360, retailers can autonomously resolve issues in a single interaction, increasing customer satisfaction and reducing operational overhead.

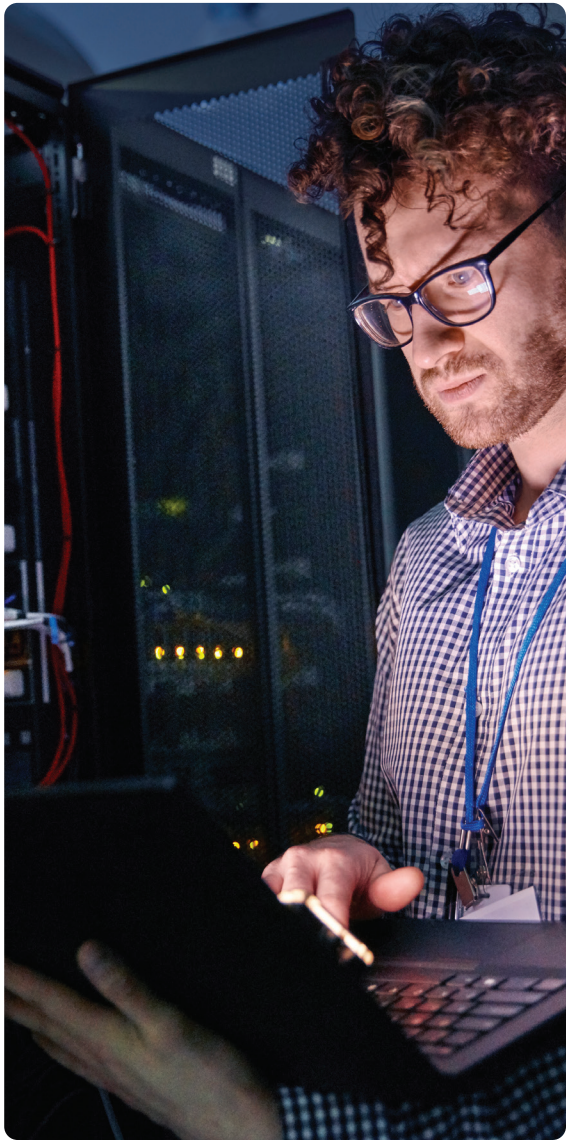

Chapter 5: The Power of Real-Time AI

Chapter 5:
The Power of Real-Time AI

As the eCommerce space continues to grow and customers are inundated with choices, the window of opportunity to influence a purchasing decision continues to shrink. Real-time delivery of AI-powered insights is more critical than ever to converting customers when their purchasing intent is at its peak. Real-time personalization requires two key capabilities: streaming data and AI at the edge.

Streaming Data

Most traditional retail data architectures rely heavily on batch processing, which moves data in large chunks, often overnight. Batch processing creates an inherent latency tax, where insights are already delayed by several hours to a full day. By the time a retailer can finally observe and respond to buying signals, it might be too late.

Streaming ingestion, processing, and analysis leveraging Apache Kafka, Apache Flink, and Apache NiFi (with MiNiFi at the edge) enables retailers to see, understand, and react to customer interactions in real time, and capture opportunities they might previously have missed.

AI at the Edge

For most retailers, it's not enough to simply implement streaming capabilities. While cloud computing is an essential component of AI development and training, for AI inference running in the cloud, the physical distance between the server and the customer is still a barrier to real-time personalization. Additionally, for retailers deploying digital, AI experiences in a physical store, fault tolerance is often necessary to prevent disruptions.

For these reasons, retailers should consider running AI inference at the edge, as close to the customer as possible. AI inference at the edge reduces latency, ensures connectivity, and keeps sensitive data local, so personalized customer experiences never disappoint.

In order to achieve real-time personalization, retailers need a specific set of technologies that enable streaming data and AI at the edge:

- A complete data in motion set of tools, including Apache Kafka, Flink, and NiFi (and MiNiFi), that enable streaming ingestion, processing, and analysis.
- The ability to deploy AI in any cloud or data center to reduce latency and enable fault tolerance.

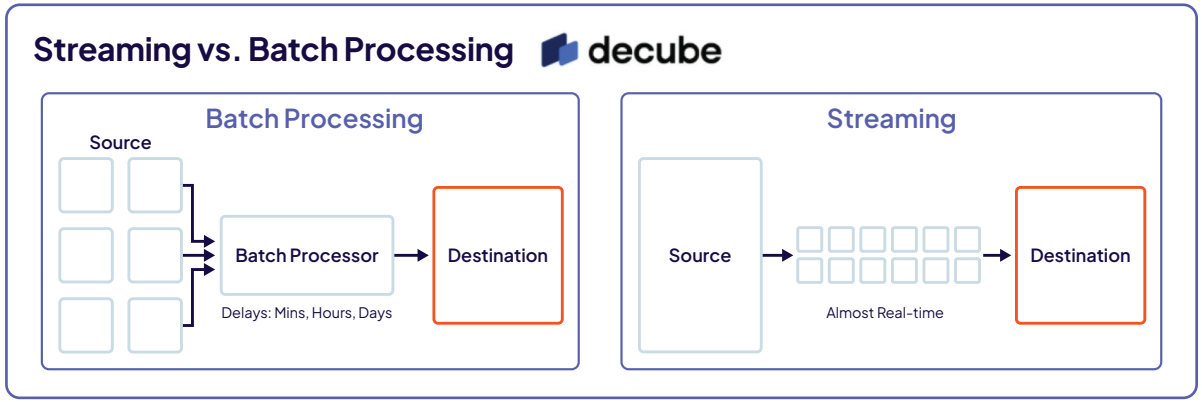

CLOUDERA

Chapter 6:
Cloudera Provides a Platform for Real-Time Personalization in Retail

Chapter 6: Cloudera Provides a Platform for Real-Time Personalization in Retail

Cloudera provides the comprehensive data architecture and AI lifecycle capabilities necessary for retailers to deliver real-time personalization at scale. By unifying siloed data and delivering actionable insights, Cloudera enables a true Customer 360 that serves as the foundation for AI-powered customer experiences at scale.

Build a Trusted Foundation with an Open Data Lakehouse. Cloudera's open data lakehouse enables retailers to integrate customer data from distributed sources, including transaction history, sentiment, and IoT data, to provide a consistent view across every department and channel.

- **Powered by Apache Iceberg**. Cloudera leverages an open table format to ensure data is never locked into a proprietary system, and enables multiple tools and engines to work on the same data simultaneously.
- **Unified Data Fabric.** Cloudera provides a single layer for governance, security, and access and discovery, ensuring consistent compliance and data quality across the entire estate, from on-premises data centers to the cloud.

Scale Intelligence with Cloudera AI. To move from descriptive to predictive analytics, retailers need to operationalize AI. Cloudera simplifies the deployment of GenAI and agentic AI and ensures AI has the customer context to make the best decisions autonomously.

- **Agent Studio.** Build and deploy agents in minutes, and employ multi-agent orchestrators to break down complex goals into smaller tasks, such as checking inventory, verifying loyalty status, and processing exchanges without human intervention.
- **RAG Studio.** Quickly build Retrieval-Augmented Generation (RAG) workflows that ground GenAI in your Customer 360 data. RAG ensures AI assistants provide accurate, context-aware answers and reduce hallucinations.
- **Fine-Tune Studio.** Tailor Large Language Models (LLMs) to your specific needs. Whether it's specialized fashion terminology or complex supply chain logic, Fine-Tune Studio enables you to customize models on your proprietary data securely.

Deploy AI Wherever your Data Lives. Develop, train, and deploy GenAI and agentic AI trained on a 360 degree view of the customer in any cloud or data center to enable personalization at the edge. Empower employees with access to contextualized customer insights using a natural language interface.

Enable Real-Time Personalization with Data in Motion. Responding to buying signals in the moment is more critical than ever for increasing conversion, loyalty, and share of wallet. Cloudera Data-in-Motion eliminates latency associated with traditional batch processing and enables retailers to act in real time.

- **Streaming Data with Kafka and NiFi.** Ingest and move data from the edge to ensure a continuous flow of real-time signals from customers.
- **Process with Flink.** Use Flink for immediate stream processing, enabling real-time delivery of personalized offers based on customer intent.

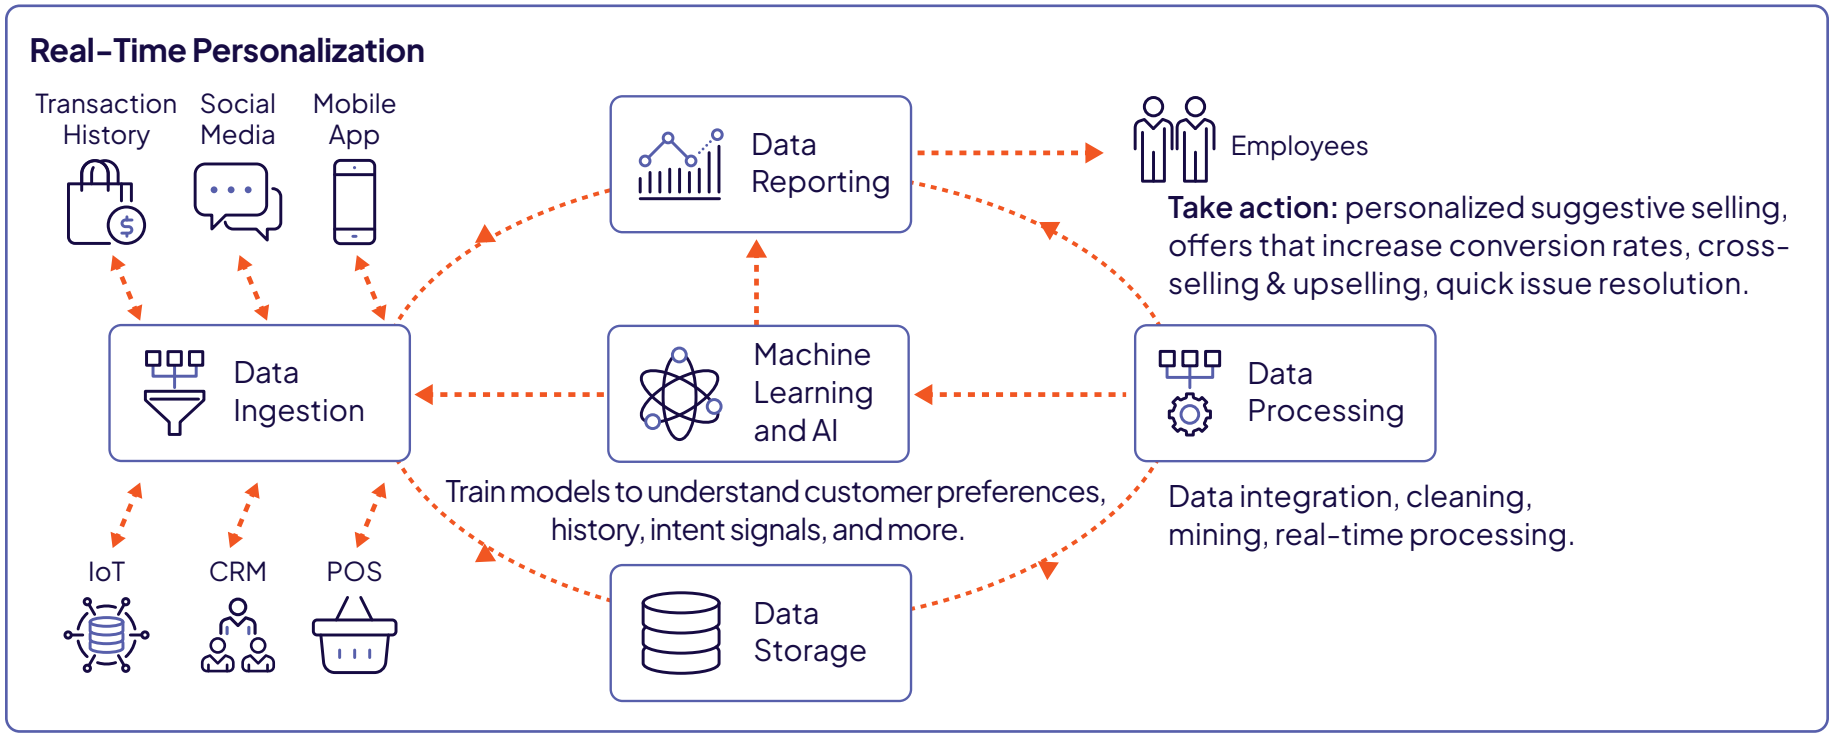
Real-Time Personalization
Transaction History
Social Media
Mobile App
IoT
CRM
POS
Data Ingestion
Data Reporting
Machine Learning and AI
Data Storage
Data Processing
Employees
Take action: personalized suggestive selling, offers that increase conversion rates, cross-selling & upselling, quick issue resolution.
Train models to understand customer preferences, history, intent signals, and more.
Data integration, cleaning, mining, real-time processing.

CLOUDERA

Ready to Get Started?

To stay competitive in today's retail landscape, companies must transition from a fragmented, stale view of the customer to a unified, real-time Customer 360. Using an open data lakehouse architecture to unify siloed customer data, data in motion to deliver real-time insights, and a low-code platform for GenAI and agentic AI development and deployment, retailers can identify and act on customer intent signals as they happen and deliver personalized, automated customer experiences at scale. Cloudera provides the hybrid platform necessary to operationalize AI in any environment, enabling retailers to capture revenue, grow loyalty and share of wallet, and ultimately make customer data a competitive advantage.

About Cloudera

Cloudera is the only hybrid data and AI platform company that large organizations trust to bring AI to their data anywhere it lives. Unlike other providers, Cloudera delivers a consistent cloud experience that converges public clouds, on-prem data centers, and the edge, leveraging a proven open-source foundation. As the pioneer in big data, Cloudera empowers businesses to apply AI and assert control over 100% of their data, in all forms, improving security, governance, and real-time and predictive insights. The world's largest brands across all industries rely on Cloudera to transform decision-making and ultimately boost bottom lines, safeguard against threats, and save lives.

To learn more, visit **Cloudera.com** and follow us on **LinkedIn** and **X**.

CLOUDERA

Cloudera, Inc. | 6220 America Center Dr, 5th Fl, San Jose, CA 95002 USA | cloudera.com

9073-002 April 28, 2026