


Developing an effective data strategy for AI in CRM

Leverage the power of Salesforce to drive better insights and efficiency in sales and service

September 2024



In today's hypercompetitive, fast-changing markets, enterprises in every industry need effective ways to preserve and enhance customer loyalty and lifetime value.

Every customer and every sale matter to the bottom line. But every customer is different—bringing unique needs, habits and expectations to each business transaction and interaction.

Customer relationship management (CRM) systems like those offered by Salesforce were created to help businesses better understand and serve their customers. Today, the predictive and generative capabilities of artificial intelligence (AI) have supercharged that potential—making it possible to personalize messages, offers and experiences at an individual level, in real time, at enterprise scale.

KEY TAKEAWAYS



AI significantly enhances CRM systems by providing deeper customer insights, automating routine tasks, and enabling personalized interactions at scale.



A well-defined data strategy is crucial for the effectiveness of AI in CRM.



Salesforce Data Cloud can help solve key requirements of your AI data strategy by leveraging prebuilt models and native connectivity, enabling faster time to value.

Rich, high-quality, accessible and actionable data is foundational to enabling trustworthy AI in CRM (and beyond). Therein lies the challenge—and opportunity—for many enterprises today.

With vast and varied datasets scattered across multiple systems (not to mention the growing amount and quality of available third-party and public data), IT and business leaders alike find themselves facing critical decisions about how to store, access, connect, enrich and activate data in ways that drive meaningful business results.

In this paper we examine key considerations and provide suggestions for how to develop an effective data strategy for enabling AI capabilities in CRM platforms.



Artificial intelligence for real customer relationships

Artificial intelligence use cases for CRM can broadly be divided into **predictive** and **generative** categories. Both are generally familiar to business leaders today—though the most appropriate solution for a specific use case might not always be intuitive.

As its name implies, **predictive AI** helps enhance decision-making and set priorities by anticipating future opportunities, events, trends or behaviors. **Generative AI (GenAI)** plays an important role in identifying new insights from large datasets and creating new assets (emails, service scripts, etc.) that support efficient interactions with customers. Both depend on historical and real-time data—from the business, from customers and from external sources—to achieve those benefits.

When it comes to the sales, commerce and service enablement capabilities of your CRM, these two types of AI can produce complementary benefits.

POTENTIAL BENEFITS OF PREDICTIVE AI IN CRM

- Enhanced customer insights
- Improved sales forecasting
- Proactive customer service

POTENTIAL BENEFITS OF GENERATIVE AI IN CRM

- More-personalized communication
- Rapid content creation
- Better customer segmentation
- Automation of routine tasks

Activating AI's potential in CRM



Implementing predictive AI and GenAI capabilities within CRM systems involves both overlapping and distinct challenges. Understanding and addressing the requirements is crucial for developing an effective data strategy.

For both types of AI, it is helpful to break down requirements into two stages: model training (merging disparate datasets into a format suitable for analysis) and inference (analyzing data to produce desired outputs).

MODEL TRAINING

Predictive AI model training necessarily involves a wide range of first-party data—including individual customer histories as well as business data about products, pricing and inventory. External datasets also play a vital role. For example, integrating medical publication data into health care CRM systems can boost pharmaceutical sales performance by alerting representatives to new research and interests of health care professionals. This enables tailored, informed, customer-centric interactions that align with the professionals' current research, enhancing relationships and ethically influencing prescribing behaviors.

Accessing and merging such a wide range of data often entails significant and time-consuming engineering efforts to bring disparate datasets into a format suitable for model ingestion.

GenAI model training is typically addressed through the use of pretrained models developed by leading AI service providers. These models enable organizations to introduce or enhance functionalities within CRM systems without the need to invest significant resources into model training from scratch. Companies can integrate these advanced capabilities swiftly, tapping into the powerful tools provided by tech vendors with minimal upfront investment in training—thereby accelerating the adoption of GenAI-driven innovations in the CRM.

Although some enterprises may opt to fine-tune pretrained GenAI models using their proprietary datasets to better align with specific business needs, this practice is relatively uncommon in the CRM realm. That's because the retrieval augmented generation (RAG) technique—whereby user interactions with a large language model (LLM) are modified by giving preference to a specific company or external dataset—can to provide high-quality outputs without the need for complex model training.

INFERENCE

Predictive AI inference tends to be perceived as relatively less complex than model training. Once a model is trained and set up, the inferencing application sends a request to the AI model's address to get predictions, using the knowledge the model gained during its training.

GenAI inference complexity can vary depending on needs. Without business-specific context, GenAI responses to inference prompts are generated based on generalized internet knowledge—which, while

adequate for certain scenarios, falls short in addressing specific, nuanced enterprise needs. Enabling GenAI models to access enterprise-specific data and information is therefore crucial. Methods such as the RAG technique discussed earlier are one key to success. Adept search capabilities are also needed. By implementing vector databases, organizations can enable semantic searching, marking a significant advancement from traditional keyword-based searches and laying the foundation for more sophisticated data retrieval methods, including knowledge graphs.





Centralized or federated data?

How about a combination?

Faced with the complexity of organizing, managing and activating data for AI in CRM, enterprises often struggle to choose between consolidating data into a centralized data lake with a single source of truth or employing federated approaches that allow data to remain dispersed across multiple systems.

The option of consolidating data, while appearing more straightforward at first glance, may require significant effort to centralize diverse datasets. Initially, organizations are required to migrate data into the central repository, a process that can be both time-consuming and costly. Then, as new sources and types of data continue to become available, core data architectures may need to evolve—creating additional work and expense.

On the other hand, fully federated data management brings its own challenges. Under this approach, data pipelines, infrastructure and applications can require significant

integration and constant monitoring to ensure reliable access to consistent, trustworthy data for AI model training and inference.

Data privacy and access regulations, corporate governance policies, and limitations imposed by third-party data providers can further influence decisions about how and where data must be stored and accessed.

Today's technologies offer a viable pathway to blend and employ *both* strategies,

empowering businesses to choose an effective solution based on their specific use cases, requirements and operational dynamics.

For example, consider the challenge of using GenAI to produce a sales email that includes relevant product and pricing information requested by a customer, while reflecting an individual sales representative's unique style and previous interactions with the customer. To expedite this process and enhance efficiency, it has become increasingly common

for CRM platforms to implement vector-oriented databases designed for semantic searches alongside traditional keyword-based search methods. This fusion of search technologies raises questions about the necessity of centralizing certain types of data—such as historical emails sent by the sales rep to prospects and detailed information about a given product's benefits for specific industries. However, internal documents and policies, which are often stored in less-structured formats, often benefit from centralization in a data lake to improve accessibility.

For use cases such as this, a hybrid hub-and-spoke data architecture—connecting both centralized and federated data systems—can be effective and efficient for activating AI capabilities in CRM. Meantime, other use cases may be served best by a purely centralized or purely federated data approach. In crafting a data strategy for your business, it's important to enable each of these approaches.



Predicting costs, generating value

Developing a data strategy for AI in CRM comes with nuanced cost considerations that must be weighed against the benefits of various approaches.

DATA INTEGRATION AND MANAGEMENT

Consolidating customer and business data from dispersed and/or legacy systems into an enterprise cloud data lake can provide a quick cost takeout as old infrastructure is retired; but if existing inconsistencies or inefficiencies are carried over in the process, this approach can lead to untrustworthy AI as well as cost overruns as data teams work to track down and address problems. Federated approaches to data management often result in higher long-term infrastructure costs as well as talent and operational costs related to maintaining and integrating data stored in multiple platforms. A federated hub-and-spoke architecture can simplify the deployment of AI use cases and shorten their time to market, but it still involves substantial initial investment in infrastructure and ongoing maintenance costs. Your data strategy for AI in CRM should include a careful cost-benefit analysis of these differing approaches.


An additional factor is the cost of valuable external data. Enterprises should encourage AI engineers and data scientists to work with business teams to evaluate and select third-party data that significantly boosts the quality and accuracy of AI models, ensuring that the investment in external data translates into tangible business value.

MODEL TRAINING

Complex data engineering efforts are often required to prepare and merge disparate datasets from various sources such as sales, marketing and customer service to create a cohesive and comprehensive data environment. Depending on the quantity, quality and diversity of data, these tasks often necessitate significant human and technological resources, driving up costs.

MODEL DEPLOYMENT

For enterprises that have already invested in advanced analytics platforms and/or large language models, the cost of integrating those capabilities with CRM platforms might be perceived as more expensive than keeping them separate. In reality, the exact cost difference may range widely—from a few basis points to magnitudes more expensive. Moreover, those additional costs should be weighed against the value that can be achieved through rapid delivery of tangible benefits, improved user and customer experiences, and the streamlining of initial investments.



Connecting it all with Salesforce

The Salesforce ecosystem provides a range of flexible options for large enterprises to leverage prebuilt data models and machine learning models to accelerate AI deployment while providing data scientists with the ability to tap into custom data science and machine learning (ML) models. As such, Salesforce can provide many of the core capabilities that fulfill your data strategy for AI in CRM, potentially shortening the path to value.

1

DATA INTEGRATION

Salesforce Data Cloud enables a 360-degree view of customer data, including sales, commerce, service and marketing data; and offers native connectivity to Salesforce applications, analytics, segmentation engines and AI/ML engines. Additionally, Data Cloud's "zero-ETL" feature enables access to data stored in enterprise cloud data platforms and other third-party systems and sources, with limited physical copying. This is a potent feature as it enables the use of a logical federated data concept. Irrespective of whether your data resides in Data Cloud or another enterprise data platform, you'll have access to that data for your analytics, AI use cases and segmentation. As a result, these capabilities can lead to higher conversion rates, enhanced customer satisfaction and stronger, more profitable relationships over time.

2

PREDICTIVE MODELING

Salesforce's Einstein Discovery can generate actionable insights, predictions and recommendations based on historical data patterns and customer behavior. It offers low-code or no-code solutions for generating models, making it easier for organizations to deploy AI-driven insights within their CRM systems. Data scientists also have the option to integrate their own predictive modeling engines through Salesforce's Bring Your Own Model capabilities, thereby extending and enriching insights with additional tools. These capabilities can help prioritize sales leads, anticipate customer issues and solutions, optimize customer interactions and reduce churn—thereby enhancing sales effectiveness and customer service satisfaction.

3

GENERATIVE AI INTEGRATIONS

Salesforce Data Cloud enables the integration of first-party and external data into GenAI prompts processed through Salesforce's Einstein Trust Layer, which helps protect the privacy and security of data while improving AI outputs. For those not utilizing the Einstein Trust Layer, Data Cloud offers an alternative approach, referred to as Bring Your Own Large Language Model, that involves creating a prompt within Salesforce using its data and leveraging an application programming interface as a stitching engine. This engine performs targeted searches on emails and product catalogs, guided by the initial Salesforce-based prompt. Adding this external information to the prompt allows it to be sent to a large language model of choice. Data Cloud also incorporates vector databases, which are particularly adept at managing document type information due to their indexing and setup capabilities. These capabilities can help automate routine tasks and provide deep insights into customer behaviors and trends, leading to more-personalized and efficient customer interactions.

4

COST EFFICIENCY AND SPEED TO VALUE

Salesforce's architecture facilitates quick access to AI-driven insights, reducing upfront capital and labor costs associated with AI integrations. By leveraging prebuilt models and native connectivity, organizations can achieve faster time to value, making the initial financial outlay more palatable. The simple integration of advanced tools and platforms developed by other AI leaders further streamlines operations, allowing Salesforce users to access predictive insights and generate personalized, customer-focused outputs directly within the Salesforce environment.





Toward an effective data strategy for AI in CRM

A generation ago, CRM platforms ushered in a new era of customer centricity in the sales and service organizations of large enterprises. Today, AI is transforming the potential of CRM systems by enabling deeper customer insights, highly personalized customer engagement, and routine task automation at scale.

The effectiveness of AI in CRM hinges on a well-defined data strategy that ensures the right data is accessible, secure, and actionable. Salesforce Data Cloud can serve as a powerful solution at the core of that strategy. It can not only facilitate quick access to AI-driven insights but also reduce the upfront capital and labor costs typically associated with AI integrations. By utilizing prebuilt models and native connectivity, organizations that implement Data Cloud can achieve faster time to value.

AS YOU WORK TO DEVELOP A DATA STRATEGY FOR AI IN CRM, HERE ARE SOME QUESTIONS TO ASK YOURSELF.

1. What datasets, internal or external, are essential for our specific AI use cases?
2. How shall we store, access and connect our diverse datasets to ensure they are actionable and trustworthy for AI applications in our CRM?
3. Should we consolidate our data into a centralized data lake, employ a federated architecture, or take a hybrid hub-and-spoke approach—and what are the cost-benefit implications of each?
4. What are the initial and ongoing costs associated with integrating AI into our CRM?
5. What processes, guardrails and key performance indicators should be implemented in our strategy to ensure optimal results and reduce risk from predictive and Generative AI?
6. How can we balance the costs of model training and inference with the expected business value and return on investment from AI-driven insights and automation?





GET IN TOUCH

David Geisinger

Managing Director
Deloitte Consulting LLP
dgeisinger@deloitte.com

Jatin Dave

Managing Director, AI & Data
Deloitte Consulting LLP
jddave@deloitte.com

Jitendra Kumar

Senior Manager, AI & Data
Deloitte Consulting LLP
jitenkumar@deloitte.com

About Deloitte

As used in this document, “Deloitte” means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Deloitte provides industry-leading audit and assurance, tax and related services, consulting, financial advisory, and risk advisory services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets, enable clients to transform and thrive, and lead the way toward a stronger economy, a more equitable society, and a sustainable world. Building on its 175-plus year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte’s approximately 457,000 people worldwide make an impact that matters at www.deloitte.com.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.