

ORACLE

Data and AI: A CIO's Guide to Success

—
It's time for an organizationwide strategy to improve productivity and unleash creativity



Table of contents

Why lead an organizationwide AI data strategy?	3
Provide an AI-ready IT infrastructure	5
Five components of an AI-ready IT stack	7
Quick win: Use AI built into your business applications	9
Quick win: Make all your data and documents searchable.....	10
Build your AI dream	12
Consider these data and AI basics	14
Free your developers to lean into generative AI.....	15
Why you want Oracle as a partner.....	16



Why lead an organizationwide AI data strategy?

Artificial intelligence sparks imaginations. New iterations of popular large language models (LLMs) add dexterity and polish to human-machine interactions, and leaders across your organization are asking, “What if...?”

CIOs can uncover beneficial AI use cases—and avoid veering toward shadow IT—by providing a unified data platform that fosters safer, higher-performing AI implementations that benefit the entire company.

With an AI-ready data platform, artificial intelligence adds urgency to the effort to derive value from your company’s operational data and domain expertise. AI makes this possible by combining natural language communications with AI models that deeply understand your organization’s data, dramatically lowering the barrier for people across your organization to query and use your vault of business and product knowledge.

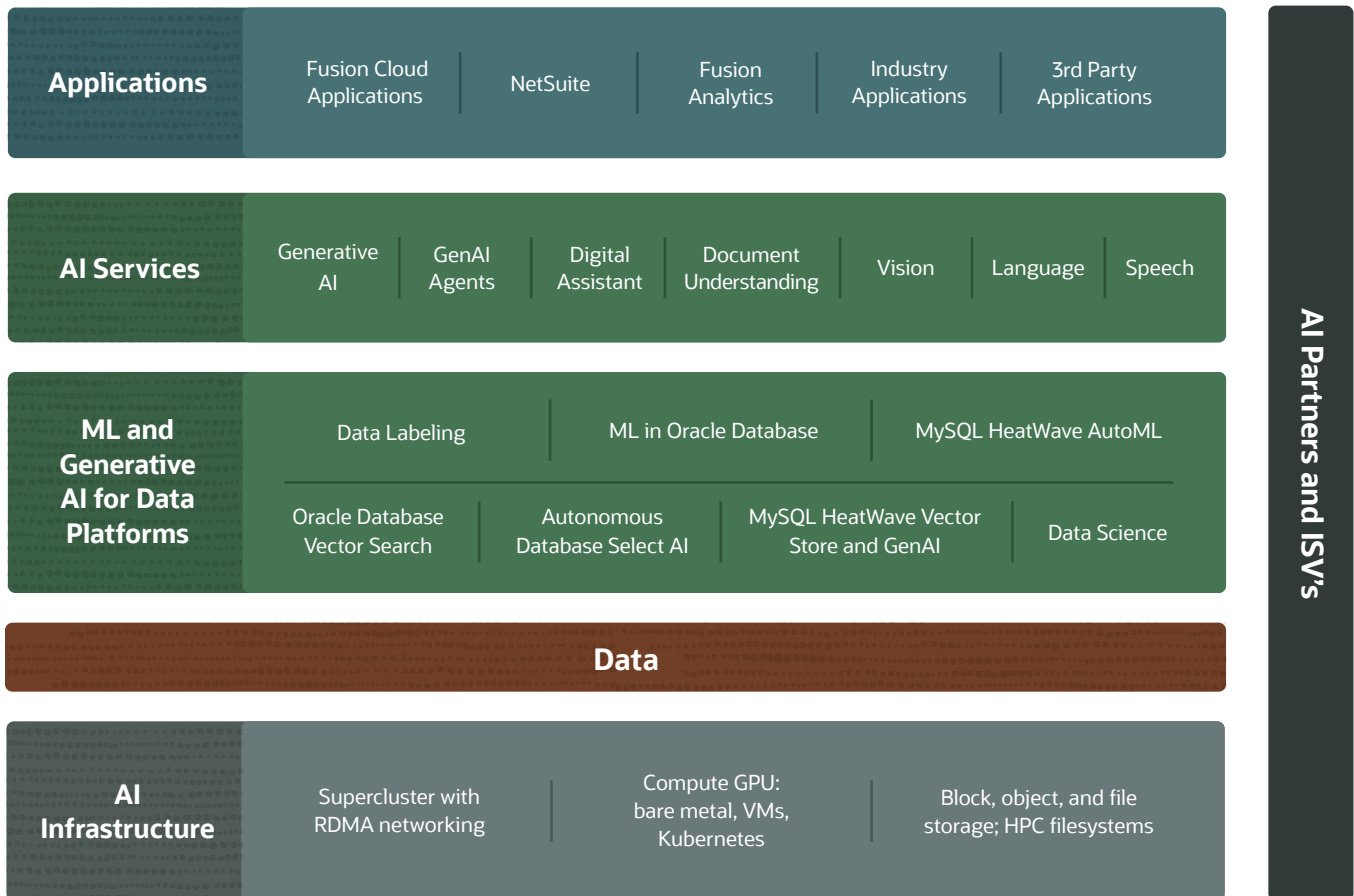
If your IT team and data analytics experts have done foundational work in your data warehouses, data lakes, and data lakehouses and implemented access controls and data governance rules, you’ve already taken the first steps. An AI-ready data platform includes centralized technology and business processes to identify, fine-tune, and deploy sophisticated AI models that take advantage of all this data.

While consumer LLMs and generative AI applications are stoking the imaginations of leaders everywhere, your job as CIO is clear: Put in place the compute power, skill sets, and companywide dataflows to tap the true promise of AI for your organization.

The Oracle AI stack

The AI stack uses cloud services to deliver a wide array of AI applications that make good use of your data.

The Oracle AI stack encompasses a range of components, integrated with Oracle’s cloud infrastructure and data management platforms to help businesses build, deploy, and manage AI-driven applications and services. The stack includes robust security and privacy features that help organizations protect sensitive data and meet regulatory requirements.



Provide an AI-ready IT infrastructure

Business leaders are asking how sophisticated AI models can improve their teams' productivity, efficiency, and creativity. CIOs can support them in discovering answers with a clear path from idea spark to production-ready.

The first step is to customize select offerings from the field of available open source and proprietary AI models. Most organizations will refine models, shaping them to their needs. The expense and technical difficulty of creating new generative AI models from scratch will be left to tech firms steeped in this cutting-edge specialty.

Fortunately, the market for adaptable generative AI models is growing and maturing. It's getting easier every day to find a model that you can fine-tune and augment to align with your data and business strategy. In the right hands, an off-the-shelf model can become a customized AI powerhouse that only your organization could create.

More targeted machine learning (ML) models are also proliferating, offering easy access to computer vision and speech, anomaly detection, forecasting, and other ML capabilities. Your internal AI program should make it easy for developers to apply ML models to applications and business operations and let them reuse models, data sets, and data labels across services.

AI terms to know



Artificial intelligence is a broad term for machines that perform tasks that were once purely the domain of humans. It includes many subtechniques, such as machine learning, deep learning, and generative AI.



Machine learning (ML) models are the basic mathematical models that allow computers to learn to accomplish tasks without specific instructions. ML powers services such as anomaly detection, forecasting, and speech recognition.

The right data science platform will help you roll out AI capabilities. It will provide a place where teams of data scientists identify, deploy, and manage ML models and LLMs—and keep them healthy with ML operations capabilities, such as automated pipelines, model deployments, and model monitoring.

Start with lower-effort implementations and grow from there. Use cases include adding AI to your current business applications and adopting AI-powered tools that make it easier for employees to access and use your storehouse of knowledge. There are also industry-specific strategies. Consider [a grocery chain that uses AI](#) to dynamically price products, personalize promotions, and go beyond self-checkout with handheld scanners that detect the items in customers' carts, then charge their accounts when they leave the store.



Deep learning is an ML subset that's carried out on artificial neural networks, which are many-layered networks modeled on the way human brains work. Deep learning is a core technique behind the abilities of large language models (LLMs) and more complex AI systems.



Generative AI can create new content. It combines many AI techniques, including ML models, neural networks, and deep learning, to understand and mimic human language, art, and decision-making. LLMs are an example of generative AI.

Five components of an AI-ready IT stack

Once the [underlying AI infrastructure](#) is in place, CIOs can clear a path for success by helping employees make the best use of data. How?

- 1 Keep it simple.** First, understand what AI capabilities are already available in your organization. Some providers of enterprise business applications are weaving helpful AI-based functions into their applications. Take advantage of those. In addition, some database providers offer AI and ML inside the database, bringing AI capabilities to where your data lives and greatly simplifying the architecture needed. These are quick wins because they're already familiar tools for employees.
- 2 Create a place to collaborate on AI models:** Make sure your data science team has a central platform to collaborate on AI workflows. Software or cloud-based platforms can guide teams through the steps, from data preparation and cleansing to model training, deployment, and inferencing, while giving technologists and businesspeople a way to access, compare, evaluate, and improve models over time. Such a platform will draw from your data warehouses, data lakes, and data lakehouses, maintaining all the data governance and access roles applied to those repositories.
- 3 Build a model library:** AI models come in many sizes and complexity levels, including LLMs and simpler ML models for services such as text analysis, computer vision, and anomaly detection. Just as developers rely on frameworks and libraries when building applications, your developers and data science teams need a repository of proven machine learning and deep learning algorithms and customized versions of LLMs. This repository often will be a component of your AI workflow platform, which allows teams to create plugins for these models and connect them to databases and other back-end systems.

4 Supercharge LLMs with your documents and data. The technology around AI and the enterprise is evolving rapidly. Two new technologies that can make LLM responses more accurate and relevant to your organization are [retrieval-augmented generation](#) (RAG) and [vector search](#). Together, they enrich an LLM with your most up-to-date documents and operational data to guide and inform AI outputs. With these technologies, a manufacturer might use generative AI to give employees and partners a new level of access to a deep catalog of product information, which they can query and explore with text or voice prompts. Similarly, a retailer might open its warehouse and CRM data to an LLM to help customer service reps understand each customer’s case and solve it quickly.

5 Facilitate streaming data for real-time AI inference. Fitting AI into your daily operations will often call for streaming data processing for AI inference. AI inference is when an AI model that has been trained for a specific task by your data science team using large amounts of curated data is moved to production, where it begins to infer, or extrapolate, conclusions from data that’s new to it. Many of the most useful AI tools for daily operations, including shipping logistics, customer service, and more, use streaming data processing so that AI outputs are based on up-to-the-minute data. Depending on the use case, the data for streaming AI inference will come from your business applications, technology and website logs, location-based data, and many other potential sources.

What is training data?

The term “training data” refers to labeled or annotated data sets used to train a machine learning algorithm. These are often very large data sets built by combining an organization’s smaller, but highly relevant, data sets with prelabeled and annotated data from available online libraries.

Complete, well-formatted training data plays a crucial role in the development and performance of AI models. Your organization will most likely choose to customize pretrained AI models and leave the time-consuming, from-scratch training process to well-funded experts in that specialty.

Again, the key is to build the unique AI tool or service that only your organization, with your data and domain expertise, can produce.



Quick win: Use the AI built into your business applications

Oracle and some other providers of enterprise SaaS applications, including ERP, HR, CRM, and supply chain management, have already built and embedded AI models for common use cases. You can take advantage of these ready-made AI implementations—and study how they were made.

- **In your ERP application**, how is AI improving efficiency and reducing human error in core business processes?
- **In your CRM application**, what criteria is AI using to identify high-potential leads and guiding sales reps with recommended actions?
- **In HR systems**, how is AI reducing time-to-hire by automating steps or improving the employee experience by helping with tasks such as performance reviews and career development planning?

Taking advantage of AI capabilities included in your business software is an easy way for employees to understand and trust AI's ability to help them do their jobs better.



Quick win: Make all your data and documents searchable

Through consumer-facing LLMs and LLM-backed search engines, people are getting used to AI that derives answers from broad swaths of online data. Now a growing list of companies, including SoundHound and Cohere, offer a secure way to provide a similar natural language interface to enterprise users. As a result, employees and partners can ask questions pertaining to operational data, such as your deep archive of documents and online communications. They can seek answers to the types of questions once reserved for analysts with carefully curated dashboards or developers proficient in programming languages.

Here's how you can securely combine your domain-specific data with the power of generative AI.

First, provide your data to your LLMs

As mentioned, RAG and vector search help enrich the output of an LLM with current organizational data, without going through the costly process of modifying the underlying model. Information can come from your structured databases and unstructured documents, such as PDFs, web pages and blogs, Word documents, news feeds, call transcripts, and online communications. If data is important to your company's performance, include it.

These data sources will be stored in a knowledge repository, such as a vector database, for fast retrieval, and then vectors will be added for use by the AI model. A vector is an easily searchable string of numbers that represent what's in a piece of content, such as the people in a photo or the topics of a text document. When someone makes a query, vectors are searched to find relevant information in the database, providing context related to the query. This contextual information, along with the query, is sent to the large language model, which uses the context to create a timely, accurate, and relevant response.


Then increase confidence in AI outputs by citing sources

Vector databases and knowledge repositories carry metadata about the sources of each piece of information, allowing the LLM to cite the source of its results and letting you easily and quickly remove information that provides unwanted or inaccurate conclusions.



How your organization implements LLMs as the face of your data will depend on your use cases.

Salespeople might ask the LLM about the most successful sales plays for a particular situation as well as details about your company's recent dealings with upsell prospects. A customer service agent might ask the LLM to explore product documentation to pinpoint a solution. The list goes on. In fact, people throughout your organization are dreaming up use cases as you read this.



Build your AI dream team

AI talent is in high demand, especially in fields such as healthcare, finance, and manufacturing. It may take a variety of technical, organizational, and interpersonal skills to bring generative AI to your company, so engage with HR and work to assemble a team with:

- ✔ **Knowledge of key programming languages:** These include Python, R, Java, and C++. A combination of these will be needed to build, customize, and implement models and to access external APIs.
- ✔ **Database expertise:** A data platform for AI requires data modeling, warehousing, and integration.
- ✔ **Machine learning expertise:** Many AI models use ML algorithms to process and analyze large amounts of data, identify patterns, and make intelligent decisions or predictions based on the patterns they observe. ML experts from your data analytics team will work to customize or augment generative AI models.
- ✔ **Knowledge of intelligent user interfaces, including voice:** Design and implement user experiences that let people intuitively interact with generative AI models in their daily workflows using text or voice prompts without coding.



What is multimodal AI?

Multimodal AI uses many data types in tandem to gain a nuanced perception of its environment. It's akin to how humans perceive the world using our five senses. Combining data types, including visual, audio, and text, helps the AI give accurate, contextualized results for an array of practical applications, such as image captioning, text-to-image generation, and even emotion recognition. Think of how an autonomous vehicle combines image recognition and road sign reading to navigate a city street.

For the data management and data science teams, this means customizing AI using video, audio, speech, images, text, and the numerical data sets from your business applications. As a result, multimodal AI training infrastructure needs special transformations and separate architectures to input, encode, combine, and eventually output meaningful predictions.

Beyond these hard skills, you need committed partners. To keep AI adoption growing, ensure your executive sponsors empower the team to quickly identify and solve problems that arise with your generative AI models and make sure the effort is properly resourced. Enthusiastic evangelists, data specialists, and product managers will help people work across data management and AI functions.

The field of AI modeling and generative AI implementation is constantly evolving. Therefore, beyond sourcing and hiring the technical experts needed for a particular function, you'll want to establish a culture of learning and create a change management team to make sure everyone understands their roles in the new AI-driven world.

Consider these data and AI basics

Is your data AI-ready? That is, is it properly cataloged and labeled, and can you track its lineage? If so, your data scientists can use it for AI modeling. Labeling is also vitally important for data privacy and security in AI outputs.

Here are additional questions to consider as you develop your strategy.

- **Is your data integrated across business functions?** Data is the foundation for AI and must be collected and stored consistently across many corporate functions to deliver meaningful benefits. The tendency for departments, including sales, marketing, customer service, and procurement, is to silo data for their own use. This approach might—or might not—have made sense before AI, but now it will stunt the effectiveness of your AI models.
- **Is your data governance in order?** Does every data domain have an owner? Is there an automated reporting structure for proving your compliance to governing agencies?
- **Where does the data come from?** To fuel AI model utility, many organizations supplement their internal data with third-party data from online marketplaces, partners, and public agencies offering useful information about location, weather, demographics, and more. How will you track the appropriateness, accuracy, and usefulness of AI outputs based on all this data?
- **Where is your data going?** If your data is going out to a managed service for a third-party AI model, how do you assess and offset the risk of unintentional leakage in that AI's output? Even internally, how do you protect against exposure of sensitive information?

Free your developers to lean into generative AI

Developers are eager to use AI in their work. Choose a data-management strategy that gives them the low-code tools to quickly innovate. New platforms are taking prompt engineering to new heights, further simplifying the way data professionals, developers, and users interact with data. With the right platform, developers can use natural language to generate features, applications, or SQL queries without writing code. This vastly improves productivity by letting developers declare their intended outcomes and get first drafts of an application instead of hand coding.



Why you want Oracle as a partner

Using data to drive business success is at Oracle's core. Now, as generative AI opens the door to gain more business value from an organization's data, that focus is even more relevant.

Oracle provides the tools, compute power, and expertise to help CIOs and their data management teams build and execute a companywide AI strategy—one that provides near-term wins and long-term innovation.

Oracle offers a full stack of cloud services, including a suite of cloud applications, databases, and a next-generation cloud infrastructure. This unique combination means [the latest generative AI](#) abilities, can be embedded in the business applications that companies use every day. In fact, Oracle has integrated AI capabilities into Oracle Fusion Cloud Applications, Oracle Database, and MySQL HeatWave. Harness Oracle's latest AI innovations to increase productivity, automate end-to-end business processes, improve decision making, and reduce the cost of doing business.

When it's time to build new AI-driven services around your data, the Oracle Cloud Infrastructure (OCI) [data platform](#) provides a secure place to collect, curate, and manage all your transactional and warehouse data and benefit from the speed of in-database machine learning. That includes all types of data, including real-time, streaming, relational, JSON, document, spatial, graph, unstructured, and open source data feeds.

Whether you're looking for an on-premises, hybrid, regulated, or public cloud solution, Oracle can support your data and AI strategy requirements.

[Learn more](#)

Copyright © 2024, Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

