ORACLE

Harness AI and Application Innovation for Success

Improve efficiency, lower costs, and drive innovation

The opportunity: Leverage AI designed for the enterprise to enhance customer experiences, improve decision-making, automate business processes, and build new products. The challenge: To implement AI cost-effectively while focusing on meaningful—and tangible—business outcomes. Oracle can help.

AI has the potential to transform every business

AI applications can help transform your business by interpreting data, recognizing patterns, making predictions, and providing intelligent responses—all while continuously learning and improving over time.

AI-powered applications can crunch lots of information quickly and accurately. That allows them to dig into customer data to provide personalized recommendations and use predictive analytics to help businesses understand buyer behavior. They can also help automate processes, optimize resource allocation, fine-tune pricing strategies, identify potential upsell and cross-sell opportunities, and lots more.

And you may be able to leverage your existing hardware, software, and personnel to adopt the latest AI-powered apps swiftly and affordably.

5 reasons to add AI to your business applications

1 Deliver excellent customer experiences

Developers will be able to build new capabilities to enhance customer experiences using generative AI to create ultraeffective chatbots and frontline assistants that can provide instantaneous responses.

AI-powered communications are more familiar—and less frustrating—because generative AI is continuously getting better at comprehension and will dispense more relevant, accurate, and complete answers over time.

Benefits: Customers get the right information at the right time. Customer representatives are free to tackle complex cases.

2 Spot trends and improve decision-making

Good use of data analysis drives business success, and AI takes it to the next level. With faster and better predictive capabilities, AI-powered applications can provide much more contextual relevance to numbers and graphs.

Benefits: Make better decisions faster, thanks to more accurate insights derived from your own data as well as a variety of external inputs.

3 Reduce costs and improve efficiency

In many organizations, employees spend much of their time performing tasks manually, such as summarizing documents and analyzing data. That not only leads to errors, but it also makes it near impossible to spot trends and anomalies given the volume of activities or data businesses handle today.

AI-based services trained with your data can help automate many tasks and allow your organization to run more efficiently.

Benefits: By leveraging cloud-based AI software and services, you can benefit from constantly evolving AI capabilities without needing to invest in new and expensive infrastructure.

4 Build new products and services

AI can help your organization create and deploy innovative applications much faster, allowing you to capitalize on recent trends to create new product categories—even entirely new lines of business—or grow your company in ways that were previously unattainable.

Benefits: Generative AI, machine learning, and vector databases provide a foundation to support innovation while a wide range of ERP, CRM, and HCM systems infused with AI capabilities can help you do more with the software you already use.

5 Build better applications faster

AI provides developers with powerful tools and capabilities that can help them produce better applications. Generative AI allows developers to create apps that use natural language processing, predictive analysis, and machine learning to understand and respond to customer needs in a more intelligent and sophisticated manner.

Benefits: In addition to facilitating the development of more-advanced apps, AI can also help with testing and optimization—for example, by helping developers fine-tune processes and identify and fix bugs before they affect users.

Drive innovation with AI and ML

Oracle's integrated suite of AI and ML services and applications gives IT leaders the power to bring innovations to market faster and make their organizations more responsive for more choice, performance, and security.

Developers can customize models using proprietary knowledge bases by using Oracle Cloud Infrastructure (OCI) AI Services, without compromising application development timelines. OCI is available at the same low price in all cloud regions and offers complete cloud services for any workload worldwide, across clouds, or in your data center.

With a strong focus on scalability, cost, performance, security, and governance for organizations of all sizes, OCI AI Services is a collection of services that enables enterprisewide collaboration and governance by allowing teams to share models, data sets, and labels across services. OCI Generative AI provides prebuilt AI language models that make it easier for developers to apply AI to their applications and business processes. And with OCI Generative AI Agents, customers can extend Cohere and Meta models with their enterprise knowledge bases and leverage retrieval-augmented generation (RAG) to gain insights from their private business data.

Take advantage of generative AI embedded into Oracle Fusion Cloud Applications, which span 18 industries and every area of the business and include Oracle Fusion Cloud Enterprise Resource Planning (ERP), Enterprise Performance Management (EPM), Human Capital Management (HCM), Customer Experience (CX), Supply Chain & Manufacturing (SCM), and more. Deploy one capability at a time or several at once from a single integrated suite with a common data model across Oracle Cloud Infrastructure and Oracle Fusion Cloud Applications to quickly meet business needs and easily tap into new capabilities as requirements change.

How Oracle helps you innovate

Oracle AI is a family of accelerated infrastructure, artificial intelligence, and machine learning services. For AI training and inferencing, Oracle's AI infrastructure offers ultralow latencies for standalone graphics processing units (GPU) and clusters with thousands of nodes.

Using AI services, developers—even those without data science expertise—can add prebuilt models to applications and operations using familiar-looking APIs and SDKs. With ML services, data scientists can build, train, and deploy models with their favorite open source frameworks or benefit from the speed of in-database machine learning.

Optimize your organization with Oracle's AI-infused apps

- Improve performance across both front- and back-office business functions.
- Discover challenges and take action, often in real-time.
- Accelerate workflows.
- Get enhanced efficiency and value delivered with the governance and security you expect from Oracle.

Generative AI can improve productivity

Oracle's approach is to help customers maximize innovations in advanced language comprehension by making it available wherever it's needed, allowing them to ideate and build their own applications incorporating generative AI.

Oracle Cloud Infrastructure Generative AI is a fully managed service available via API that allows customers to seamlessly integrate versatile language models into a wide range of use cases, including writing assistance, summarization, and conversational chat. Oracle offers a choice of proprietary Cohere models or Meta's open source Llama 2 model. Use the prebuilt models or customize them with your data, allowing models to respond to requests with up-to-date and trusted information from your source systems through retrieval-augmented generation.

Build on your existing applications faster with AI

With our experience building enterprise applications, integrating AI tools into your existing software platforms becomes seamless. These tools allow you to enhance business systems with conversational interfaces, provide users with access to information without the need for specialized training, automate common business processes and tasks, and more, all of which can save you time and money.

Accelerate LLM training and inference

Compared with traditional compute, GPUs significantly accelerate training and inference workloads when implementing AI solutions. OCI's AI infrastructure is designed to provide you with access to cutting-edge GPUs in an elastic, scalable, and affordable way.

One of OCI's key differentiators is a simple, flat Ethernet network that supports remote direct memory access (RDMA), a network technology that lets multiple servers in a cluster access shared memory without bottlenecks. Because of OCI's superior cluster network bandwidth, you can scale to incredibly large clusters compared with other hyperscale cloud providers. Even at scale, the system's low latency makes OCI's AI infrastructure incredibly fast.

Why is OCI so fast? The infrastructure is designed and built around the OCI Supercluster technology right from the start, making OCI's high performance computing and AI/ML infrastructure rival on-premises infrastructure for less cost.

"The limitless opportunities for AI-driven innovation are helping transform virtually every business. NVIDIA's collaboration with OCI puts the extraordinary supercomputing performance of NVIDIA's accelerated computing platform within reach of every enterprise."

Manuvir Das
Vice President of Enterprise Computing, NVIDIA

Leverage all the advantages of Oracle's AI infrastructure

- OCI often surpasses the performance of dedicated, custom on-premises compute clusters while providing the elasticity and usage-based cost model of the cloud.
- OCI Supercluster, part of OCI's AI infrastructure, is ideal for training generative AI, including conversational applications and diffusion models.
- With support for up to tens of thousands of NVIDIA GPUs, OCI Compute bare metal instances and virtual machines can be used for many applications, such as computer vision, natural language processing, recommendation systems, and generative AI.

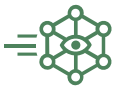

Oracle and NVIDIA have partnered to deliver infrastructure that accelerates AI workloads and the adoption of machine learning models.

Each customer can train generative AI with an OCI Supercluster that can currently scale to include tens of thousands of NVIDIA GPUs.

Leverage open source and existing technology stacks

Oracle has made significant contributions to GraalVM, Java, MySQL, and other popular platforms for application innovation. And OCI has embraced open source and multicloud architectures to offer customers maximum flexibility.

Oracle APEX AI-assisted development and Oracle Database's Multilingual Engine help low-code developers quickly build secure, scalable, and feature-rich enterprise applications by using AI to convert natural language prompts into SQL queries and then create the desired app components. The Multilingual Engine even allows developers to write server-side functions in other languages, if they choose, while benefitting from the low-code paradigm.

Want to add even more data to your chosen AI model? No problem. AI vector search within Oracle Database helps you store semantic information about documents, images, and other unstructured files as vectors, or numerical representations of text. This lets large language models quickly find closely related terms by adding generative AI capabilities to Oracle Database development tools—including APEX—so users can write applications and queries without coding. AI Vector Search in Oracle Database 23c also enables organizations to perform searches on a combination of both business and semantic data—which is what enterprises need to implement effective AI solutions—particularly when both types of data are managed by a single database.

What about data scientists? Our machine learning services help specialists collaboratively build, manage, and deploy machine learning models with their favorite open source frameworks. Many popular AI frameworks, software, and management tools are available as open source through NVIDIA's NGC portal. The enterprise version of this open source collection, called NVIDIA AI Enterprise, provides enterprise support and is available through the Oracle Cloud Marketplace.

Oracle offers industry-leading tools for managing data and AI models

Using our wide variety of computer science, data science, and developer tools, your technical teams can build and deploy cloud native apps with familiar technologies, such as Kubernetes, Jupyter, Python, R, REST APIs, and SQL.

They can also collaboratively build, train, deploy, and manage machine learning models with their favorite open source frameworks and take advantage of templates to build new AI/ML models.

OCI Data Science can accelerate building, training, deploying, and managing ML models throughout the model lifecycle by providing a comprehensive, easy-to-use suite of tools in one place.

NIBIO

Oracle APEX is the world's most popular enterprise low-code application development platform. It's used by the Norwegian Institute of Bioeconomy Research (NIBIO) to improve sustainable forest management.

Read NIBIO's story

NIBIO analyzes massive amounts of data with ML models on OCI

- The Norwegian Institute of Bioeconomy Research uses cutting-edge drones, sensors, and other machinery to gather large amounts of data about the country's forests to balance conservation efforts with the demands of logging.
- NIBIO hopes this technology will help plan the future of Norway's forests and inform decisions with impacts lasting hundreds of years.
- Using Oracle Database, APEX, and OCI Data Science, NIBIO built a system to store, analyze, and access the diverse types of data collected from the forests.

Accelerate your entry into AI with access to computational power at scale

AI requires considerable computing power to train models and apply that learning to business problems. Fortunately, with Oracle you can access cloud-based infrastructure for advanced applications such as customized generative AI and LLMs.

That means you can use enterprise-level AI without investing in massive server farms.

OCI was designed to provide the best price-performance in the industry. Further, we offer the same consistent infrastructure pricing, regardless of your region, and you may save 50% compared with other cloud providers—so you get predictable performance at a predictable price.

When it comes to technology, Oracle stands apart. For example, when training AI models, OCI's unique cloud architecture quickly shuttles data between advanced CPUs and GPUs to deliver the best possible performance. OCI Supercluster is powered by the newest NVIDIA GPUs in the market, connected with an ultralow-latency network. This provides the speed needed to quickly train generative AI models at scale.

OCI Supercluster includes GPU-powered bare metal instances. Combined with ultralow-latency networking and high performance storage, these instances offer AI companies an ultrafast platform for machine learning, image processing, model training, inference computation, physics-based modeling and simulation, and massively parallel HPC applications.

Each bare metal instance has eight NVIDIA GPUs, and these bare metal instances can be combined to a maximum of up to 32,768 NVIDIA GPUs. Many customers are using OCI's advanced capabilities to train and deploy large-scale AI applications.

Why choose Oracle for AI?

- OCI's AI infrastructure rivals or betters the performance of dedicated custom on-premises compute clusters while providing the elasticity and consumption-based costs of the cloud.
- OCI offers industry-leading cluster network bandwidth, which can scale to incredibly large clusters compared with other hyperscale cloud providers.
- Customers can use the Oracle Generative AI service to benefit from on-demand elasticity and scalability, predictable price-performance, and the ability to create private model endpoints.
- We have decades of experience in data management systems such as Oracle Database and MySQL. Our in-database AI and ML capabilities help eliminate complexity to deliver the highest levels of performance, scalability, and availability.
- We train specialized models that are unique to verticals and our industry-leading SaaS solutions, and you can further refine these prebuilt models using your own data, so your models understand your business like no other.

Ready to implement Oracle AI solutions?

We meet you where you are in your AI journey with a variety of embedded and managed services across our infrastructure layer, platform services, and business applications. Working with AI may seem challenging—but it's much simpler if you're working with a company that has created an entire cloud of integrated services from data to apps.

That's the Oracle AI advantage.

Learn more

Explore Oracle's AI portfolio
Request a demo
Learn more about OCI
Cloud strategies for the CIO

Connect with us

Call +1.800.ORACLE1 or visit oracle.com

Outside North America, find your local office at oracle.com/contact