

Boost AI & analytics using the software your business already counts on

How Microsoft, Oracle and SAP platforms
provide an easy on-ramp for machine and
deep learning

intel®

Contents

| | |
|---|---|
| Introduction | 2 |
| Oracle Database | 3 |
| SAP HANA | 4 |
| Microsoft SQL Server | 5 |
| A universal platform for high-performing AI workloads | 6 |
| Conclusion | 8 |
| Learn More | 8 |

Introduction

From personalizing a customer's shopping experience, to automating predictive maintenance on the factory floor, or identifying emerging industry trends and opportunities, artificial intelligence (AI) can help drive business performance and strengthen revenues. It takes us beyond the capabilities of traditional analytics (identifying past patterns to predict future events) to the ability to sense, reason and direct next actions automatically and in real-time.

IT teams are often pressured to deliver on the AI expectations and demands of the lines of business within their organizations. However, this need not mean making costly investments or implementing complex new technologies. Most modern enterprises already have much of what they need in place, through their use of enterprise software packages and platforms from major independent software vendors (ISVs) such as SAP, Oracle and Microsoft. These ISVs continue to invest in deep learning and other AI capabilities that can be adopted with minimal or no additional investment.

In this paper we'll explore how three of the most common ISV offerings are supporting or enabling AI. We will also examine how your Intel® architecture, with unique built-in AI acceleration and optimizations co-developed with the ISVs, can help you get optimal performance from them like no other.



Oracle Database

Oracle Database is a multi-modal database management system typically used to run online transaction processing (OLTP), data warehousing and mixed database workloads. Its AI capabilities tend to focus on traditional machine learning, as these types of algorithm tend to be best suited to the data and workloads it supports

The platform uses AI in two ways:

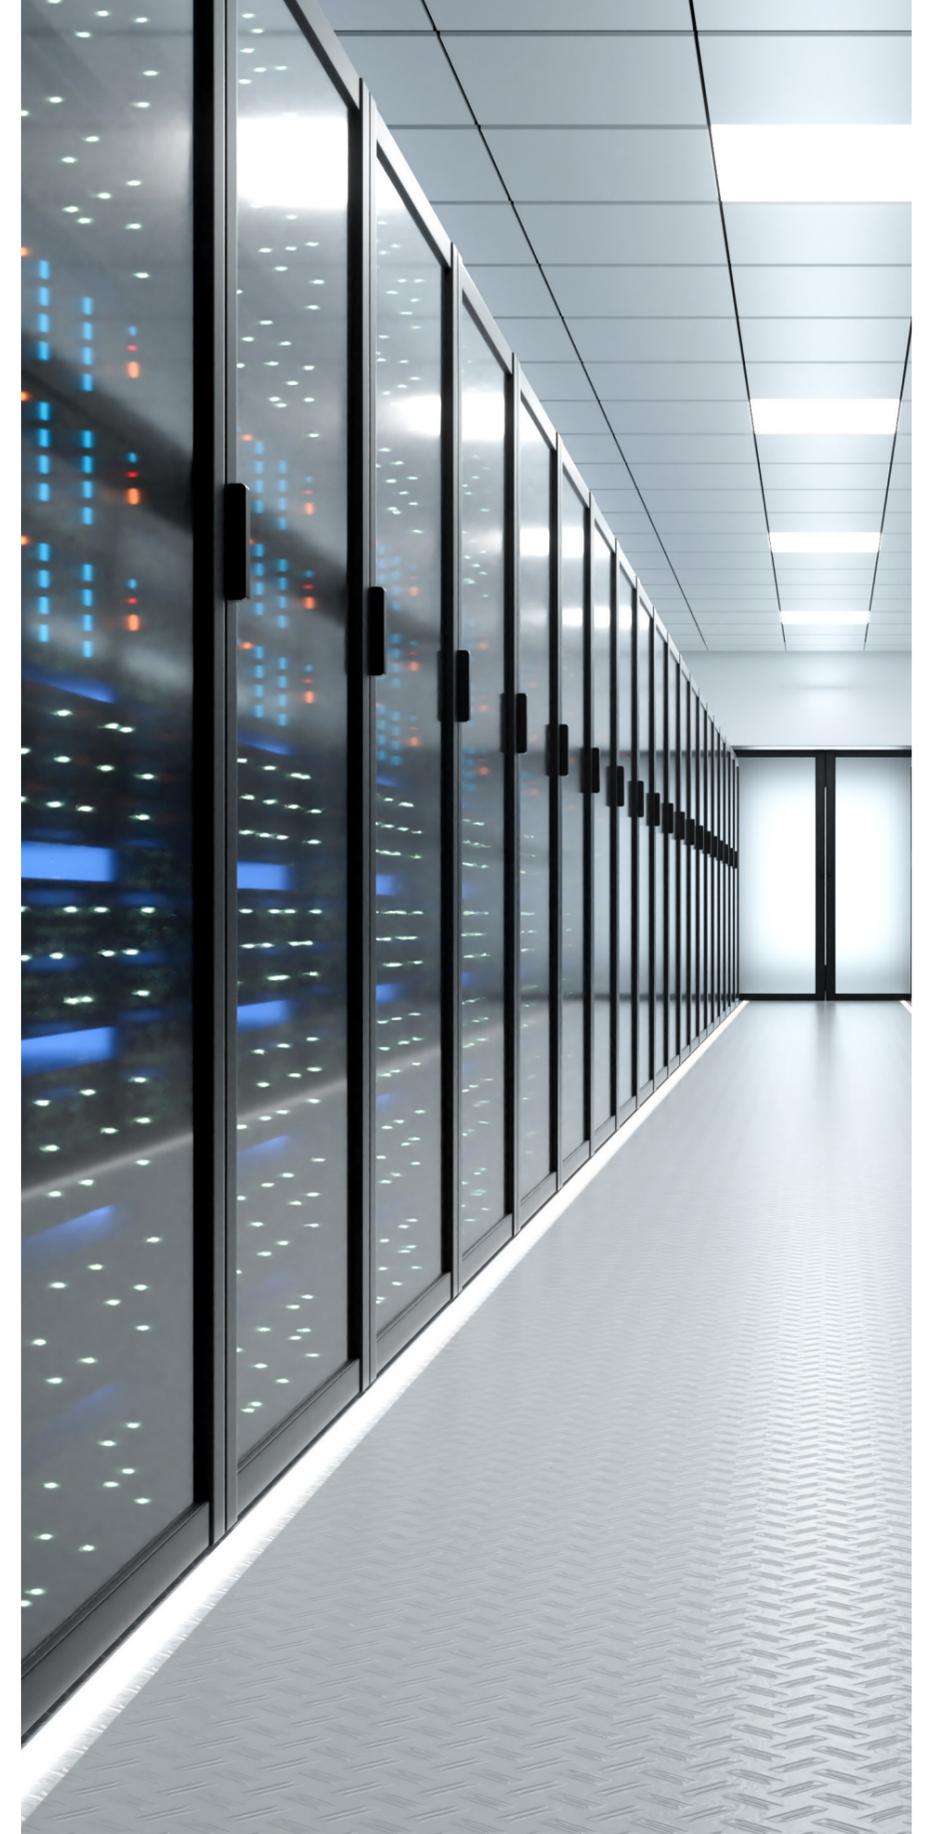
- Machine learning-based tools are embedded into Oracle applications running on the Oracle Database platform to enable business users to investigate their data. The Oracle Machine Learning module—delivered free of charge with Oracle DB 12.2 and later—also allows data science teams to build their own machine learning models in order to address the requirements of a given team or organization.
- Specific AI-driven algorithms or features are applied to certain Oracle applications in order to enhance the service or user experience. For example, a chatbot feature may be included in an existing financial management application so users can interact with the app in a more natural way, or a fraud detection app may use machine learning to enhance accuracy and automate responses.

Intel and Oracle have worked together for decades to concurrently develop new hardware and software features that work optimally together. This ongoing co-engineering approach allows for constant exploration of how Intel® technology can be used to enhance the performance of Oracle's offerings now and in the future.

For example, Oracle uses Intel technology in its own hardware as part of its Engineered Systems (ES), such as Oracle Exadata (both on-premise and in the cloud). This is the ES designed to run Oracle Database products, and so supports multiple AI workloads.

How to take advantage of AI on Oracle Database

The good news for Oracle Database customers is that these AI capabilities are already available for you to use through the Oracle Machine Learning module, whether you are using on-premise Oracle Database or Oracle Exadata instances, or accessing the capabilities through Oracle Cloud. Oracle adds new AI features and capabilities to its platform through scheduled updates to Oracle DB, which automatically become available to license holders with no need to purchase additional licenses. Encourage your line of business users to investigate the features available today in their most commonly used applications, and keep an eye out for updates over time.





SAP HANA is an in-memory database (IMDB) that supports a diverse range of workloads. Its S/4 HANA offering supports online transaction processing (OLTP) workloads including all business suite applications, while BW 4/HANA runs online analytics processing (OLAP) workloads like data warehousing.

The AI capabilities in SAP HANA are delivered using a built-in predictive analytics library (PAL), which contains a number of universal predictive algorithms that can execute directly against the data held in the database. These include algorithms for AI functions like clustering, classification, association and time series analysis.

Additional AI capabilities can be gained using the SAP Data Intelligence platform, which includes algorithms for more complex workloads, such as deep learning, conversational AI, video analytics, and robotics process automation. For these algorithms, training takes place on the SAP Data Intelligence platform, before the algorithm is ported over to run inference on the data in SAP HANA as needed.

Like Oracle, SAP HANA enables customers to apply these AI capabilities to whichever application or business use case they need. For example, an HR team may use a number of different SAP HANA applications to support their workflows. They may choose to bring in an AI-driven tool to help automate the process of reviewing new resumés, checking for key words and narrowing down the number of applications that the team needs to manually review.

How to take advantage of AI on SAP HANA

Organizations using SAP HANA today can take different steps depending on where they currently are in their AI journey. For those just starting out or interested in traditional machine learning use cases (such as OLTP), explore the built-in options available through the PAL library.

To evolve your AI capabilities further, you may choose to invest in the SAP Data Intelligence platform in order to access a broader range of AI algorithms for more complex use cases. Intel provides a large set of optimized libraries for popular AI machine learning models that enable Intel Xeon Scalable processors to accelerate your AI workloads. Before doing this, it is worth making sure your data in SAP HANA is ready to support the AI you plan to run, and that your supporting infrastructure is equipped with the Intel technologies that will help accelerate and optimize your chosen workloads.

Microsoft SQL Server

Microsoft SQL Server is a relational database management system (RDMS) that manages and serves data for use by other applications, which—like those of Oracle and SAP—can cover a wide range of industries and use cases.

The 2016 release of SQL Server introduced machine learning capabilities that allow AI and analytics to be run on individuals' machines, without having to send data across the network or be limited by the memory of their own devices. The platform includes distributions of Microsoft R and Python, which contain commonly used data science packages, as well as some proprietary AI packages.

SQL Server 2019 built in additional support for data-hungry workloads like AI by including Apache Spark and Hadoop Distributed File System (HDFS) to create a unified data platform called '[big data cluster](#)'. This supports large-scale data processing and machine learning up to exabytes, including unstructured data formats such as images and speech. Organizations can use either SQL Server Machine Learning Services in their master instance, or Spark, to run AI model training in R, Python or Java. Whichever they choose, they can use open source libraries such as TensorFlow or Caffe—which also have optimizations for Intel® technology—to train the models.

SQL Server environments run across bare metal, virtual machines (VMs) and the cloud, supported across all these by Intel hardware and software optimizations for AI.

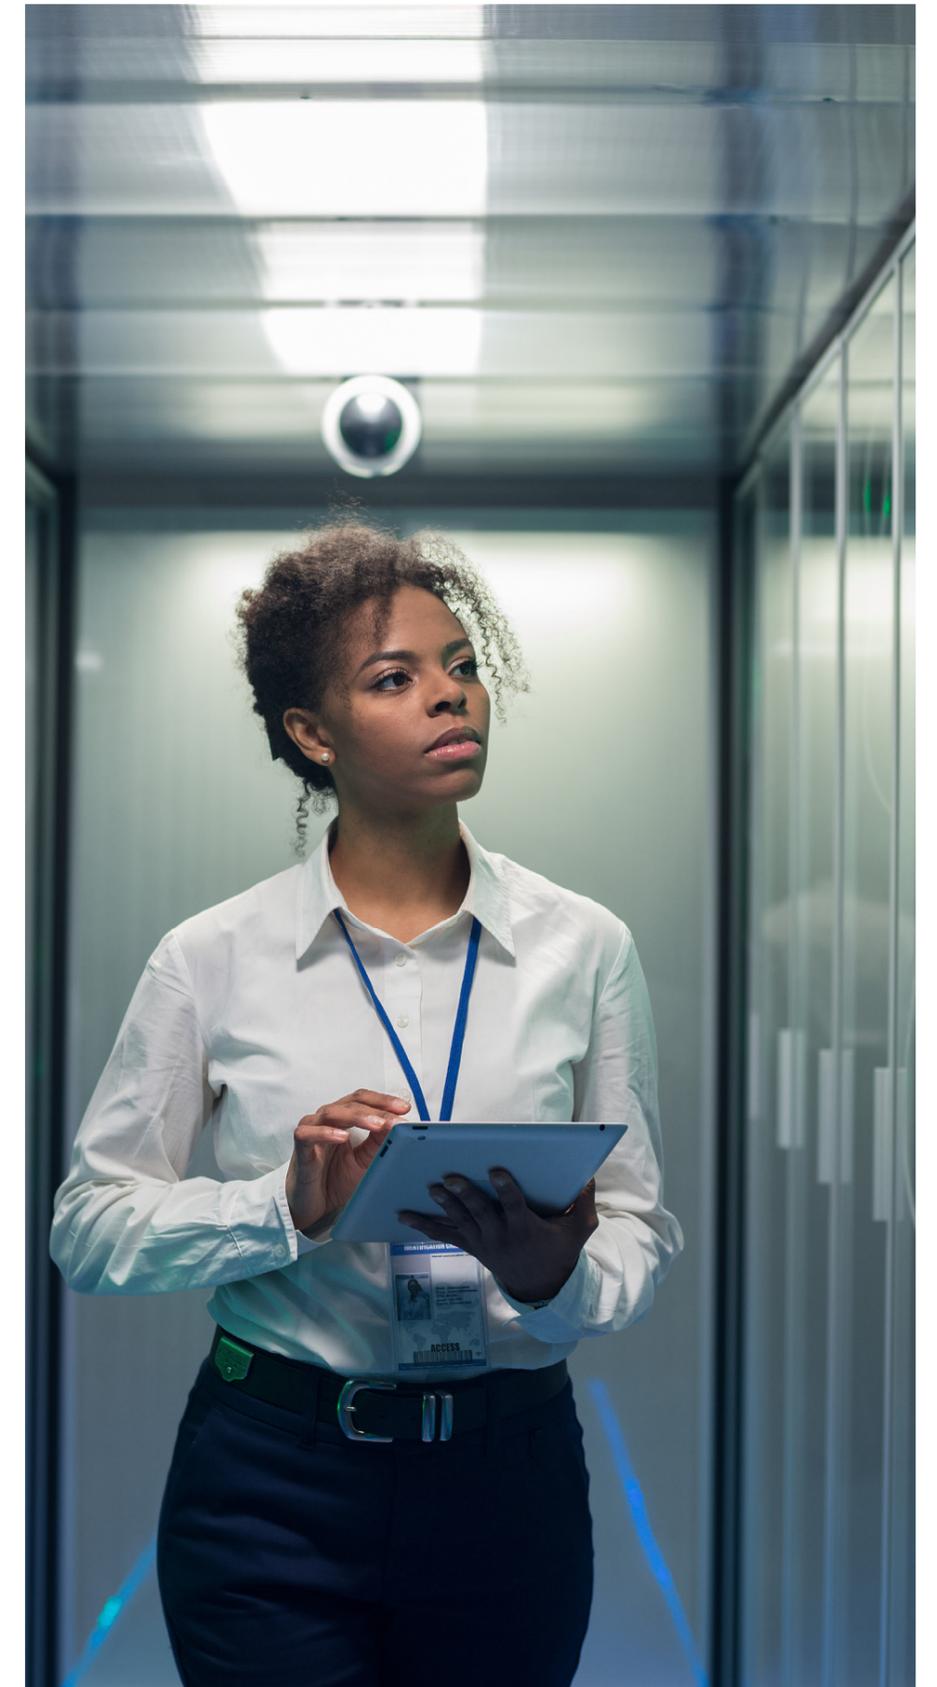
Microsoft is also a participant in the Open Neural Network Exchange (ONNX), which develops an open framework that allows different machine learning models to be used interchangeably. ONNX runs on SQL Server for the edge, or Microsoft Azure SQL, enabling full edge-to-cloud AI. It includes support for Intel DL

Boost, which helps further improve performance for workloads such as product recommendations, image recognition and convolutional neural networks (CNNs). This helps accelerate inference, enabling more data to be processed in the same time, and representing a potential cost saving for organizations that use cloud-based services billed by the hour.

How to take advantage of AI on SQL Server

Those using the SQL Server platform today have these AI capabilities available to them already, with no need to configure anything. Simply introduce your pre-trained model and use your SQL Server environment to identify the data for training and then inference.

For guidance on how to introduce these models, Microsoft provides step by step guidelines, such as these for [deploying ONNX models on SQL Edge](#).



A universal platform for high-performing AI workloads

Intel works closely with its ecosystem—including the major ISVs—to build solutions optimized for analytics and AI performance.

Processing power

- The latest 3rd generation Intel® Xeon® Scalable processors evolve Intel's 4- to 8-socket processor foundation for AI-infused, data-intensive workloads. They offer built-in AI acceleration that further supports real-time AI.
- The latest evolution of Intel® Deep Learning Boost (Intel® DL Boost) technology—an innovation of the 2nd generation Intel Xeon Scalable processor—makes the 3rd gen Intel Xeon processor the first general-purpose CPU to offer built-in bfloat16 instructions. Building on existing support for Vector Neural Network Instructions (VNNI), the addition of bfloat16 instructions makes mainstream AI training more widely deployable for workloads like image classification, speech recognition, and language modelling. These features help make these general-purpose CPUs a more cost-effective option than purchasing expensive GPUs specifically for AI workloads.

In-memory performance

- Delivered with 3rd generation Intel Xeon Scalable processors, Intel® Optane™ persistent memory (PMem) offers massive memory capacity to support data-hungry in-memory AI workloads, and [performs up to 600x faster than flash-based SSD](#)¹.

- ISVs like Microsoft, Oracle and SAP can therefore offer increased capacity for large in-memory workloads like real-time analysis. This is generally cost-prohibitive with other types of memory like DRAM.
- The technology also enables the low latency and high throughput that real-time workloads demand.
- Intel Optane PMem can also be used in virtualization environments, meaning applications can run smoothly across complex hybrid IT infrastructures.

Big data acceleration

- Intel® Advanced Vector Extensions 512 (Intel® AVX-512) offer a set of instructions to accelerate complex in-memory workloads like analytics and AI.
- Intel® Key Protection Technology (Intel® KPT), a feature of the Intel Xeon Scalable processor. Used to protect sensitive encryption keys at runtime, this feature adds an additional layer of protection for the Exadata platform.
- Intel® QuickAssist Technology (Intel® QAT), which helps accelerate compression and decompression of data so it can be processed and analyzed quickly.

Software optimizations

- The Intel® Distribution of OpenVINO™ toolkit allows organizations to implement computer vision and deep learning solutions quickly and effectively across multiple applications and hardware platforms.

- Application developers can use the Intel® oneAPI Deep Neural Network Library to improve the productivity and enhance the performance of their deep learning frameworks. This includes taking advantage of the 3rd Intel Xeon Scalable processor's built-in AI acceleration capabilities like bfloat16 support to speed up training for AI models.
- Optimizations for common languages including R, Python and Java running on Intel® architecture are available to help further accelerate AI and machine learning workloads on common ISV platforms.

| | ORACLE DB | SAP HANA | MICROSOFT SQL SERVER |
|----------------------------------|---|--|---|
| AI Features | <ul style="list-style-type: none"> • In-app ML for data analysis • Oracle Machine Learning module • AI-enabled features to enhance user experience | <ul style="list-style-type: none"> • Predictive Analytics Library (PAL) offers built-in predictive algorithms • SAP Data Intelligence Platform can add more complex AI • AI capabilities can be added to any app running on SAP HANA | <ul style="list-style-type: none"> • In-app ML for business users • Big Data Cluster allows training of AI models • ONNX open ML framework |
| Key Enabling Intel® Technologies | <ul style="list-style-type: none"> • Intel® Xeon Scalable processor with Intel® Deep Learning Boost • Intel® Optane™ persistent memory • Intel® AVX-512 • Intel® Key Protection Technology | <ul style="list-style-type: none"> • Intel® Xeon Scalable processor with Intel® Deep Learning Boost • Intel® Optane™ persistent memory • Intel® AVX-512 • Intel® QuickAssist Technology • Intel® Distribution of OpenVINO™ toolkit • Intel® oneAPI Deep Neural Network Library | <ul style="list-style-type: none"> • Intel® Xeon Scalable processor with Intel® Deep Learning Boost • Intel® Optane™ persistent memory • Intel® oneAPI Deep Neural Network Library |
| How to Use AI | <ul style="list-style-type: none"> • AI capabilities already built-in – no additional licenses needed • Encourage line of business users to explore features in their apps of choice • Look out for new features being added through updates | <ul style="list-style-type: none"> • Explore existing options through PAL • Take the next step by investing in SAP Data Intelligence Platform | <ul style="list-style-type: none"> • AI capabilities already built in – no additional licenses needed • Consult Microsoft guidance for deploying ONNX models |

Table 1: Intel® technologies enabling AI capabilities in common ISV platforms



Conclusion

By taking full advantage of the AI and machine learning capabilities already made available to you by your ISVs of choice, you can equip your business users today with many of the tools they need. As these common enterprise applications already run on and are optimized for your Intel® architecture-based environment, there's no need to purchase additional hardware or even software licenses. This creates a low-cost, low-risk opportunity to experiment with the use cases and types of AI that are most appropriate for your organization.

These applications can help you hit the ground running with data-driven decision making, with minimal need for incremental investment.



Learn More

- Webpage: [Oracle Database](#)
- Webpage: [SAP HANA](#)
- Webpage: [Microsoft SQL Server](#)
- Blog: [Oracle Machine Learning](#)

¹ Oracle evaluation. Testing conducted by Oracle in May 2020: <https://blogs.oracle.com/database/persistent-memory-primer>

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

Other names and brands may be claimed as the property of others.

0121/JL/CAT/PDF

345421-001EN