

# Oracle Database In-Memory

*Powering the Real-Time Enterprise*

**18<sup>c</sup>** ORACLE<sup>®</sup>  
Database

## KEY FEATURES

- In-Memory Column Format accelerates Analytics by orders of magnitude
- Dual-format architecture combines the best of column and row formats
- Many In-Memory optimizations, including In-Memory storage indexes, In-Memory compression, SIMD vector processing
- In-Memory Expressions and In-Memory Join Groups further increase performance
- Automatic In-Memory controls the content of the In-Memory column store
- Scale-up on the world's largest SMPs with 1000s of CPU threads
- Scale-Out on RAC clusters with in-memory fault tolerance, and the ability to isolate analytics from OLTP on different servers
- Maximum Availability Architecture including support for In-Memory column store on Active Data Guard Standby
- 100% compatible with existing applications
- SQL In Silicon - database processing built into silicon enables 10x greater performance and 2x better compression

## KEY BENEFITS

- Dramatically faster Data Warehouses, Analytics, Business Intelligence, Dashboards, Reports
- Fast ad-hoc analytics without the need to pre-create indexes
- Reduces space consumed by analytic indexes that are no longer required
- Faster mixed workload OLTP

Oracle Database In-Memory – part of Oracle Database 18c Enterprise Edition – adds in-memory database functionality to existing databases, and transparently accelerates analytics by orders of magnitude while simultaneously speeding up mixed-workload OLTP. With Oracle Database In-Memory, users get immediate answers to business questions that previously took hours.

Oracle Database In-Memory delivers leading-edge in-memory performance without the need to restrict functionality, or accept compromises, complexity and risk. Deploying Oracle Database In-Memory with any existing Oracle Database compatible application is as easy as flipping a switch - no application changes are required. Oracle Database In-Memory is fully integrated with Oracle Database's renowned scale-up, scale-out, storage tiering, availability, and security technologies making it the most industrial-strength offering on the market.

Oracle Database In-Memory is now recognized as a leader in The Forrester Wave™: In-Memory Databases, Q1 2017. The full report is available on [oracle.com](http://oracle.com).

The ability to easily perform real-time data analysis together with real-time transaction processing on all existing applications enables organizations to transform into Real-Time Enterprises that quickly make data-driven decisions, respond instantly to customer demands, and continuously optimize all key processes.

## Dual-Format Architecture

Oracle Database In-Memory optimizes analytics and mixed workload OLTP, delivering outstanding performance for transactions while simultaneously supporting real-time analytics, business intelligence, and reports.

This breakthrough capability is enabled by the “dual-format” architecture of Oracle Database In-Memory. Up to now, databases have forced users to store data in either column or row format. Column format is highly efficient for analytics, but imposes very large overheads when used in OLTP environments. Similarly, row format enables

- Concurrent analytics and OLTP
- No database size limit
- Transparently blends memory, flash, and disk to achieve highest performance and lowest cost
- Industrial strength availability and security
- Robustness and maturity of Oracle Database 18c
- Easy implementation and management

extremely fast OLTP, but is less optimized for analytics. The only way to optimize for both OLTP and analytics has been to copy data from OLTP systems to analytic systems using complex ETL processes that add a great deal of expense and latency.

The dual-format architecture of Oracle Database In-Memory eliminates this tradeoff by representing tables simultaneously using traditional row format and a new in-memory column format. The Oracle SQL Optimizer automatically routes analytic queries to the column format and OLTP queries to the row format, transparently delivering best-of-both-worlds performance. Oracle Database automatically maintains full transactional consistency between the row and the column formats, just as it maintains consistency between tables and indexes.

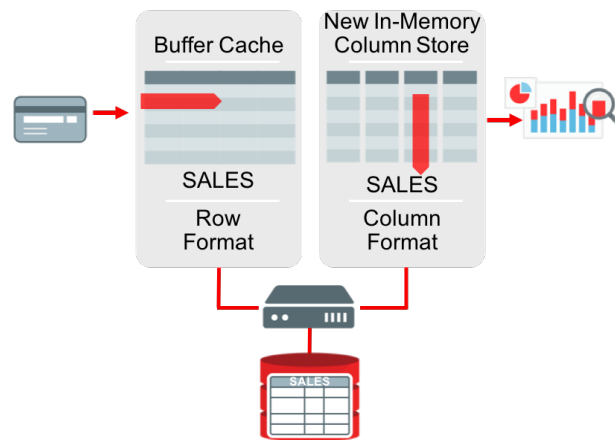


Figure 1. Oracle's unique dual-format architecture

Dual row and column formats are a unique feature of Oracle Database In-Memory.

High performance transactions in combination with column formats is a unique feature of Oracle Database In-Memory. Most databases don't allow efficient OLTP when using column formats.

*"Now we can run time-sensitive analytical queries directly against our OLTP database. This is something we wouldn't have dreamt of earlier."*

*Arup Nanda  
Enterprise Architect  
Starwood Hotels and Resorts*

The new column format is a pure in-memory format. Tables are stored on disk using Oracle's existing row-based or hybrid columnar formats. Since there is no persistent columnar storage format, there are no additional storage costs or storage synchronization issues. Changes to the purely in-memory column format are very fast because they don't need expensive persistent logging.

Having both a column and a row-based in-memory representation does not double memory requirements. Oracle uses its highly optimized buffer cache management algorithms to keep only actively accessed row data in memory. Decades of experience has shown that caching a small percentage of data blocks in memory eliminates the vast majority of storage I/Os, and flash caching eliminates virtually all the rest. Therefore most of the memory capacity in a database server can be allocated to the column format.

Oracle's in-memory column format uses sophisticated compression to expand memory capacity and improve query performance. Compression ratios vary from 2X – 20X, depending on the option chosen and redundancy in the data. The compression method may be different across columns, partitions or tables. For example, some table partitions can be optimized for scan speed, others for memory footprint, while others may be optimized to efficiently handle frequent DML operations.

## Comprehensive In-Memory Optimizations

Oracle Database In-Memory implements state-of-the-art algorithms for in-memory scans, in-memory joins, and in-memory aggregation. For example:

- Analytic queries typically reference only a small number of the columns in a table. Oracle Database In-Memory minimizes work and maximizes performance by accessing only those columns needed by a query, and by processing these columns directly without having to decompress them first.
- Tables are logically split into sections, and minimum and maximum values of every column are maintained for every section of a table. This allows queries to quickly skip table sections that only contain data outside of the range of data needed by the query.
- Some columns have many repeated values. For example a column storing the geographic region of each sales transaction will have many repetitions of the same regions. Oracle Database In-Memory compresses these repeated values to save memory, and optimizes processing by executing query predicates just once for each unique column value (e.g. once per unique region).
- Modern microprocessors support SIMD (Single Instruction for Multiple Data) vector processing instructions to accelerate graphics and scientific computing. Oracle Database In-Memory can use these SIMD vector instructions to process multiple column values in a single CPU clock cycle.

*“Oracle Database In-memory is a game changer for OLTP, DW, and mixed workloads. It dramatically improves the performance of all types of analytical queries.”*

*Liviu Horn  
AVP Database Management  
McKesson Health Solution*

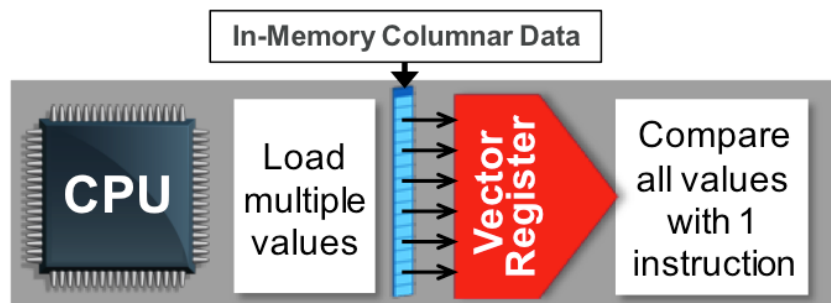


Figure 2. SIMD vector processing scans billions of values per second

- In-Memory table joins take advantage of the new columnar compressed format by converting join conditions into filters applied during very fast data scans or by executing the join on the compressed values within the join columns.
- *In-Memory Join Groups*, if declared, can further speed up in-memory hash joins by eliminating the overhead of decompressing and hashing column values. This further reduces CPU cycles and improves join performance.
- *In-Memory Expressions* enable frequently evaluated expressions to be materialized into the In-Memory column store. Analytic workloads typically spend a considerable amount of time on expression evaluation for each row returned by a query. Once in the column store, all in-memory optimizations seamlessly apply to In-Memory Expressions: vector processing, storage index pruning, etc. In-Memory Expressions are also maintained automatically as the underlying table columns are updated.
- In-Memory aggregation algorithms leverage the column format to speed up analytic queries and reports that aggregate large amounts of data.

Together, these optimizations enable Oracle Database In-Memory to run queries at the astounding rate of billions of rows per second for each CPU core. Analytics that

Most other databases incur large overhead when data in column format is updated as part of mixed OLTP workloads.

*“Oracle Database In-Memory option will permit us to remove roughly half the indexes from large mixed-workload databases, simultaneously speeding up both complex analytical queries and OLTP transactions.”*

*Andrew Zitelli  
Principal Software Engineer with Honors  
Thales-Raytheon Systems*

Isolation of analytic workloads on separate servers of a RAC cluster is a unique feature of Oracle Database In-Memory.

Many other in-memory databases severely limit database size, and require costly DRAM for all data including infrequently accessed data

Many other in-memory databases have limited scale-up capabilities and no scale-out or very immature scale-out implementations.

previously took hours or days to run now completes in seconds, enabling real-time business decisions.

## Mixed Workload OLTP

Mixed workload databases run reports and ad hoc queries in addition to OLTP transactions. It is common to create dozens of indexes on important tables purely to speed up reports. The Oracle Database In-Memory column format eliminates the need for most of these analytic indexes by delivering performance similar to having an index on every column, but with much less transaction overhead. Removing analytic indexes speeds up OLTP operations since these analytic indexes no longer need to be maintained by every transaction. Further, eliminating the need for analytic indexes removes the time-consuming tuning and administration required to identify and create them. Most importantly, users can now issue any query they want and achieve excellent response times. Users are not limited to running queries or reports that have been pre-optimized using analytic indexes.

Oracle Real Application Clusters (RAC) can be combined with Oracle Database In-Memory to isolate analytic workloads from OLTP workloads by running them on different servers of a cluster.

## Cost Effective In-Memory Processing for any Database Size

Oracle Database In-Memory does not require all database data to fit in memory. Users can choose to keep only performance sensitive tables or partitions in memory. Less performance sensitive data can reside on much lower cost flash or disk.




Figure 3. Data may reside in memory, flash or disk with no size limits

Queries execute transparently on data residing on all three tiers - memory, flash and disk - enabling Oracle Database In-Memory to be used with databases of any size. If *Automatic In-Memory* is enabled, the contents of the In-Memory column store can be automatically managed based on heat map statistics.

Oracle Database In-Memory is fully compatible with Oracle’s Multitenant database architecture, allowing consolidated databases to take advantage of a combination of fast in-memory and low-cost storage technologies.

Scale-out for real-world OLTP workloads with column formats is a unique feature of Oracle Database In-Memory.

*“Full support for RAC scale-out means Oracle Database In-Memory can be used on our largest Data Warehouse, enabling more near real-time analytics.”*

Sudhi Vijayakumar  
Senior Oracle DBA  
Yahoo Inc

*“Downtime is extremely costly for our business. Oracle’s In-Memory architecture takes the right approach to balancing real-time speed with continuous availability.”*

Jens-Christian Pokolm  
Analyst IT-DB Architecture & Engineering  
Postbank Systems AG

## Scale-Out and Scale-Up

Oracle Database has been optimized and tuned for decades to scale-out on clusters of servers and scale-up on SMP servers. Oracle can parallelize a single SQL statement across thousands of processor cores, or run many concurrent users simultaneously with full isolation and consistency.

Oracle Database In-Memory builds on these technologies to scale-out to very high memory and CPU capacities by using all of the memory and processors in a cluster of servers. Unlike traditional in-memory databases, Oracle Database In-Memory has fully transparent and highly efficient scale-out capabilities for both analytics and OLTP. For example, the Exadata Database Machine X7-8 can be configured with up to 18 racks with up to 3 8-socket database servers per rack yielding a total of 54 database servers with 972 terabytes of DRAM and 10,368 processor cores. Oracle Database In-Memory automatically distributes tables across the in-memory column stores of all the instances in a cluster. This distribution is similar to the striping of data in a storage subsystem. In-memory SQL execution transparently queries and combines data across all the instances of the cluster using parallel execution processes. Oracle Database In-Memory further optimizes scale-out query processing by co-locating frequently joined partitions on the same instance to enable local partition-wise joins. Inter-node communication on Engineered Systems uses Oracle Database’s ultra-fast InfiniBand Direct-to-Wire protocol to achieve both very low latency and high throughput.

In addition to being able to scale-out, Oracle Database In-Memory can also scale-up to very high memory and CPU capacities on large SMP servers with terabytes of memory. For example, the SPARC® based Oracle M8-16 SuperCluster Machine can be configured with up to 16 terabytes of DRAM and 16x32-core SPARC® M8 processors in a full rack. The Intel® Xeon based Oracle Exadata Database Machine X7-8 can be configured with 18 terabytes of DRAM and 576 processor cores per rack. Large SMP servers are well suited for in-memory workloads because all memory is accessible to all processors over an extremely high-speed memory network. Oracle Database In-Memory optimizes performance on large SMP servers by preferentially scheduling threads on each CPU to access in-memory data that is located on DRAM local to that CPU.

## Industrial-Strength Availability and Security

Oracle has worked for decades to build, refine, and harden its high availability and security technologies. Oracle Database In-Memory inherits all the proven functionality of Oracle Database, including the sophisticated and robust high availability solutions embodied in Oracle’s popular *Maximum Availability Architecture* (MAA).

Because the Oracle Database In-Memory column format resides purely in-memory, it does not change storage formats for data files, redo logs, or undo. Therefore, Oracle’s renowned backup, recovery, disaster recovery, and replication technologies work transparently to the in-memory column format with no changes to functionality, operations, or administration. Oracle’s industry leading security technologies also operate completely transparently with Oracle Database In-Memory. Complete support and compatibility with Oracle’s industrial strength availability and security eliminates

the need to accept new business risks in order to achieve leading edge in-memory performance.

## In-Memory Fault-Tolerance

Most other in-memory databases cannot maintain duplicate in-memory copies and therefore will stall processing for many minutes or hours while data is reloaded from storage onto another node

When a server node fails, the in-memory data on that node is lost. Queries can continue to run on surviving nodes, but it takes time to repopulate the in-memory data from storage, and during this time analytic queries will run much slower.

The fault-tolerance feature of Oracle Database In-Memory eliminates this slowdown on Oracle Engineered Systems by optionally duplicating data across the nodes of a cluster. Just as storage subsystems stripe and mirror data across disks to achieve high performance and high availability, Oracle Database In-Memory distributes and duplicates in-memory data across the nodes of a cluster. If a node fails, queries can transparently use the duplicate copy of data on surviving nodes.

In-memory fault-tolerance can be enabled at a table or partition level to balance the extra memory consumption of in-memory fault tolerance against its availability benefits. For example, customers may choose to enable fault tolerance for recent partitions to ensure fast response times for critical transactions while accepting slower response for queries on historical data in the event of a node failure.

## Database In-Memory on Active Data Guard

Most other in-memory databases have limited and very immature high availability and security features.

With Oracle Database, it is possible to create, populate, and maintain an In-Memory column store on a standby database in an Active Data Guard configuration. This allows full use of the CPU and memory resources on the standby for analytic workloads and allows for flexible placement of the In-Memory column store across the primary and its affiliated standby databases.

Having the same data populated into the In-Memory column store on the standby increases the fault-tolerance for the In-Memory column store: analytic workloads can continue to run even if there is temporary outage of the primary or the standby.

Alternatively, the In-Memory column store on the primary and the standby databases can be populated with different datasets, increasing the capacity of the In-Memory column store beyond what the primary database alone could accommodate.

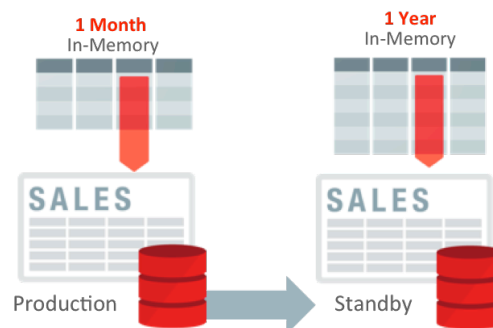


Figure 4. In-Memory column store fully usable on Standby with Active Data Guard

## Database In-Memory FastStart

In-Memory Fast-Start allows significantly faster initial population of the In-Memory column store following a database planned or unplanned outage. With In-Memory Fast-Start, the contents of the In-Memory column store can be written to storage in the form of a SecureFile LOB in the designated "FastStart tablespace" once the column store has been initially populated and automatically maintained. During a database restart, the contents of the In-Memory column store can be very quickly read from the SecureFile LOB and placed into memory, thus avoiding the need to repopulate the base table row data into compressed column vectors.

## Automatic In-Memory

In Oracle Database 18c, Automatic In-Memory is available to automatically manage the contents of the IM column store. If the sum of the space of the segments that have been enabled for in-memory exceeds the available memory in the IM column store then Automatic In-Memory will kick in and manage the IM column store space using heat map statistics. Using access tracking, column statistics, and other relevant statistics segments can be automatically evicted from the IM column store to make room for the population of more active segments.

In addition, Automatic Data Optimization (ADO) performs automated Information Lifecycle Management on database objects. ADO extends automated policy-based management to the In-Memory column store as well. It is possible to set ADO policies that:

- Mark an object as a candidate to be populated into the In-Memory column store.
- Increase the In-Memory compression level for a table as the table data cools (number of accesses reduces).
- Evict a table altogether from the In-Memory column store after a certain number of days or after a certain period of inactivity. This is ideal for system with a rolling window access patterns. Automatically evicted inactive data from the In-Memory column store makes room for new data.

Using Automatic In-Memory and ADO allows the In-Memory column store to become self-managing.

## SQL in Silicon

A few years ago, Oracle set out to build the world's fastest In-Memory database technology as well as the world's fastest microprocessor for running database workloads. The culmination of this effort is SQL in Silicon: Database In-Memory functionality built natively into the SPARC® M8 microprocessor, specifically engineered for optimal performance for Database In-Memory.

On standard microprocessors, Database In-Memory executes SIMD vector instructions in order to process multiple column values in parallel. While Vector Processing does achieve billions of rows per second scan performance, SIMD vector instructions were originally designed for graphics and are not a natural fit for database processing.

Further, SIMD vector instructions consume CPU cores just like standard processor instructions. As depicted in Figure 5, the SPARC® M8 processor has native *database accelerators* (DAX), which can directly and independently process Database In-Memory operations. Thus, instead of having to translate a scan that finds matching city names into vector operations, this scan can be directly offloaded to DAX while the CPU cores can do other work. SPARC® M8 SQL in Silicon provides 32 lightweight DAX engines for Database In-Memory using only an additional 1% in chip real estate

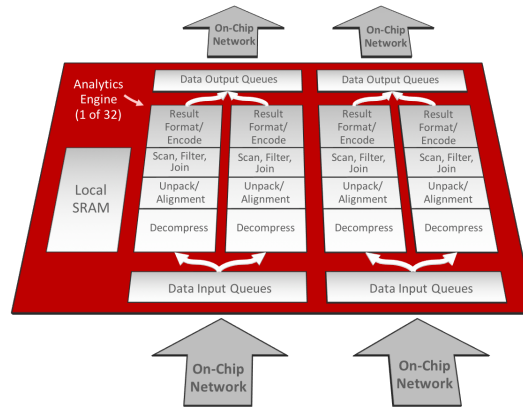


Figure 5. SPARC® M8 processor with SQL-in-Silicon Database Accelerators

SQL In Silicon provides two major benefits:

**Offload Processor Cores:** The DAX processors offload work from the standard database cores enabling much more database processing to occur in parallel.

**Increased In-Memory Capacity:** Database In-Memory allows in-memory data to be compressed for greater in-memory capacity. Fast decompression is the key to running in-memory scans efficiently. Typically, symbol-based compression, in which values are replaced by smaller fixed-size symbols (dictionary compression is an example of such a scheme) provides the fastest scan performance since data does not need to be decompressed while being scanned although it does not provide the most optimal compression factor. Database In-Memory also features higher levels of compression such as bit-pattern compression using a proprietary algorithm known as Oracle Zip. Oracle Zip is optimized for the fastest possible decompression speed but it does impose a performance penalty since decompression is needed before data can be scanned. With DAX, the Oracle Zip decompressor is also built natively into Silicon and incurs negligible performance penalty compared to symbol based compression. As a result it is possible to have **2x the in-memory capacity** with SQL in Silicon.

## Lower Costs

The ultra fast performance delivered by Oracle Database In-Memory not only improves response times, it also lowers costs and improves productivity. For example: hardware and software for both servers and storage can potentially be reduced; employees are more productive because they no longer need to wait for slow reports; and database administrators spend less time tuning.

*“Oracle Database In-Memory made our slowest financial queries faster out-of-the box; then we dropped indexes and things just got faster.”*

*Evan Goldberg  
Co-Founder, Chairman, CTO  
NetSuite Inc.*



*“With Oracle Database In-Memory, we saw tremendous improvement in our query performance, and dropping indexes reduced our database size by 80%.”*

Francois Bermond  
Data & Analytics  
Schneider Electric

*“Utilizing Oracle Database In-Memory against our JD Edwards ERP suite for real time summarization has the potential to radically change the way we deliver reporting, analysis and insights to our business leaders.”*

Michael Macrie  
Chief Information Officer  
Land O’Lakes

No other in-memory database can match the functionality of the Oracle Database, and most impose restrictions on SQL syntax or database functionality when column formats are used. These restrictions make it impossible to run most existing applications without extensive modifications or limitations.

## Real-Time Enterprise

Traditionally simple transactions execute in real-time, but answering business questions that require detailed data analysis can take hours. Oracle Database In-Memory takes advantage of massively scalable hardware with new in-memory data structures and algorithms to immediately answer any question. Oracle Database In-Memory processes data at a rate of **billions** of rows per second rather than millions. Moreover, analytics can run directly in OLTP databases, further reducing delays and improving accuracy by avoiding having to transform and copy data to a different systems.

The ability to easily combine real-time data analysis with real-time transaction processing on all existing applications enables organizations to become *Real-Time Enterprises* that:

- Make data-driven decisions based on immediate and accurate answers
- Respond instantly to customer demands for information, choices, personalization, and engagement
- Continuously optimize all key processes including sales, marketing, manufacturing, staffing, costing, etc. using detailed, up-to-date data

Using Oracle Database In-Memory, organizations can become Real-Time Enterprises that out-innovate competitors, delight customers, and improve the bottom line.

## Oracle In-Memory Applications

Oracle is embracing the benefits of Oracle Database In-Memory throughout its products. Each of Oracle’s Applications - including Oracle Fusion Applications, Oracle JD Edwards EnterpriseOne, Oracle PeopleSoft, Oracle Siebel, Oracle E-Business Suite, and Oracle Hyperion—is developing new In-Memory modules that leverage Oracle Engineered Systems and Oracle Database In-Memory to transform critical but slow business processes into real-time processes. For example:

- Cost and Profitability Analysis that took 57 hours now runs in minutes
- Financial Position Analysis that took over 4 hours now runs in seconds
- Sales Order Analysis that took days now runs in less than a second
- Consumption Driven Planning that took 13 hours now runs in minutes

Users and application developers can use Oracle Database In-Memory to make similar improvements in their own applications.

## 100% Compatible

Oracle Database In-Memory delivers leading-edge in-memory performance without the need to compromise functionality or compatibility. Oracle customers have invested hundreds of billions of dollars developing applications on top of Oracle Database. These applications make extensive use of Oracle Database’s rich SQL and PLSQL functionality, data types, optimizations, and capabilities. Oracle Database In-Memory enables all existing applications to achieve in-memory performance with near zero effort.

Oracle Database In-Memory is designed to be completely and seamlessly compatible with existing applications. No changes are required to use it with any application or tool

that runs against Oracle Database. Analytic queries are automatically routed to the column store by the SQL optimizer, and transactional semantics are guaranteed by the database.

With Oracle Database In-Memory, users do not need to rewrite, rebuild, or migrate their applications. Users can quickly adopt in-memory technologies to speed up existing applications, and continue to focus their efforts on developing new applications that improve their business.

*“In terms of how easy the in-memory option was to use, it was actually almost boring. It just worked - just turn it on and select the tables, nothing else to do.”*

*Mark Rittman  
Chief Technical Officer  
Rittman Mead*

## Easy to Implement and Manage

In addition to being compatible at the application level, Oracle Database In-Memory is easy to implement and manage. Enabling Oracle Database In-Memory is as easy as setting the size of the in-memory column store and identifying tables or partitions to bring into memory. Background processes populate data from storage into in-memory columns while the database remains fully active and accessible. Oracle Enterprise Manager makes it easy to monitor and measure the benefits of in-memory columns.

## Summary: Highest Performance, Maturity, and Compatibility

Oracle Database In-Memory transparently accelerates analytic queries by orders of magnitude enabling real-time decisions. It dramatically accelerates Data Warehouses, Data Marts, and Mixed Workload OLTP environments. Oracle Database In-Memory implements a unique dual-format architecture that delivers fast analytics together with high-performance OLTP.

Oracle Database In-Memory is easily deployed under any existing application that is compatible with Oracle Database. No application changes are required. Oracle Database In-Memory uses Oracle's mature scale-up, scale-out, and storage-tiering technologies to cost effectively run any size workload. Oracle's industry leading availability and security features all work transparently with Oracle Database In-Memory, making it the most robust offering on the market.

Extreme performance for both analytics and transactions enables organizations to continuously optimize processes and make rapid data-driven decisions thereby transforming into Real-Time Enterprises that are extremely agile and efficient.

## ORACLE DATABASE IN-MEMORY

## POWERING THE REAL-TIME ENTERPRISE


Speed Up Analytics by Orders of Magnitude	Oracle Database In-Memory transparently extends industry-leading Oracle Database 12c with columnar in-memory technology. Users get <u>immediate answers to business questions that previously took hours</u> because highly optimized in-memory column formats and SIMD vector processing enable analytics to run at a rate of <u>billions of rows per second per CPU core</u> .
Unique Architecture Runs Analytics in Real-Time while Accelerating Mixed Workload OLTP	Column format is optimal for analytics while row format is optimal for OLTP. Oracle Database In-Memory uses both formats simultaneously to allow <u>real-time analytics on both Data Warehouses and OLTP databases</u> . Indexes previously required for analytics can be dropped, accelerating mixed-workload OLTP.
Compatible with All Existing Applications	Deploying Oracle Database In-Memory with any existing Oracle Database-compatible application is as easy as flipping a switch, <u>no application changes are required</u> . All of Oracle's extensive features, data types, and APIs continue to work transparently.
Industry-Leading Scale-Up	Oracle's highly mature scale-up technologies enable application transparent In-Memory scale-up on SMP computers with up to tens of terabytes of memory and thousands of CPU threads. Data is analyzed at the enormous rate of hundreds of billions of rows per second with outstanding efficiency and no feature limitations.
Industry-Leading Scale-Out	Oracle's highly mature scale-out technologies enable application transparent In-Memory scale-out across large clusters of computers with <u>100s of terabytes of memory</u> and thousands of CPU threads. Data is analyzed at the enormous rate of <u>trillions of rows per second</u> with no feature limitations.
Industry-Leading High Availability and Security	<u>Oracle's renowned Availability and Security technologies all work transparently</u> with Oracle Database In-Memory ensuring extreme safety for mission critical applications. On Oracle Engineered Systems, In-Memory fault tolerance duplicates in-memory data across nodes enabling queries to instantly use an in-memory copy of data if a node fails.
Cost Effective for Even the Largest Database	Oracle Database In Memory <u>does not mandate that all data must fit in memory</u> . Frequently accessed data can be kept In-Memory while less active data is kept on much lower cost flash and disk.
Powering the Real-Time Enterprise	The ability to easily perform real-time data analysis together with real-time transaction processing on all existing applications enables organizations to transform into Real-Time Enterprises that quickly make data-driven decisions, respond instantly to customer demands, and continuously optimize all key processes.

## CONTACT US

For more information about [insert product name], visit [oracle.com](http://oracle.com) or call +1.800.ORACLE1 to speak to an Oracle representative.



## CONNECT WITH US



## Integrated Cloud Applications &amp; Platform Services

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. 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.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0116



Oracle is committed to developing practices and products that help protect the environment