

Oracle® for SAP® Database Update



Oracle Database for SAP Roadmap

ORACLE DATABASE 11g FOR SAP APPLICATIONS



Oracle/SAP Support Status:

- 11.2.0.4: supported until 12/2020 (12/2018*)



Supported Platforms:

- Generic Servers (On-Premise)
- Oracle Engineered Systems
- Oracle Cloud (OCI only)



New Features/Options:

- Advanced Compression (Tables, Indexes)
- Tablespace Encryption



Entry Note:

- SAP Note 1431800

ORACLE DATABASE 12c FOR SAP APPLICATIONS



Oracle/SAP Support Status:

- 12.1.0.2: supported until 07/2021 (07/2019*)
- 12.2.0.1: supported until 11/2020



Supported Platforms:

- Generic Servers (On-Premise)
- Oracle Engineered Systems
- Oracle Cloud (OCI, ExaCS)



New Features/Options:

- Oracle Database In-Memory
- Advanced Compression (ILM)
- Oracle Multitenant



Entry Notes:

- SAP Note 1914631 (12.1)
- SAP Note 2470660 (12.2)

ORACLE DATABASE 18c FOR SAP APPLICATIONS



Oracle/SAP Support Status:

- supported until Q2/2021



Supported Platforms:

- Generic Servers (On-Premise)
- Oracle Engineered Systems
- Oracle Cloud (OCI, ExaCS)



New Features/Options:

- In-Memory Optimized Arithmetic
- Polymorphic Table Functions (SAP: FOR ALL ENTRIES)



Entry Notes:

- SAP Note 2705272
- SAP Note 2660020

* Extended support fees waived by Oracle

DEAR SAP CUSTOMER,

The relationship of Oracle Corporation and SAP SE has been and is still based on our long history, a rich heritage of joint developments and a bright future – for the benefit of our mutual customers. Both companies have had an ongoing commitment to our tens of thousands of joint customers for over 25 years. Our recently renewed and longstanding reseller and support agreements provide enhanced access to Oracle Database technology, Oracle Cloud Infrastructure and world class customer support.

Oracle will support SAP Business Suite and SAP BW as long as SAP will be supporting them. With every new release we will provide latest database technology and let customers make use of more and more SAP application optimizations.

Starting 2018, new releases of the Oracle Database software are provided annually. In addition, a new numbering schema has been implemented: Instead of the traditional version number (11g, 12c), the software versions are now designated by the year of their release (18c, 19c, etc.). These annual software releases will be made available to SAP on Oracle customers as well.

An overview of the versions that are currently available or will be available in the near future can be found in the infographic on the left page. For additional details see SAP Notes 1174136 and 2606828.

The Oracle product strategy provides flexibility and choice across the whole IT infrastructure. Therefore, a majority of mid-size to the largest enterprise SAP customers in every industry entrust their application deployments to the Oracle Database.

By choosing the Oracle Database and Database options, SAP customers significantly benefit through the ongoing innovations without disruption. By taking a closer look, eight differentiators have been identified which explain in detail why the Oracle Database is the first choice for running SAP applications. The Oracle Database brings best performance and scalability, deployment flexibility, availability and reliability, support for database consolidation and very large databases, database security, manageability and self-management as well as integration of hardware and software.

All SAP customers can enjoy these benefits. It makes no difference whether you run the Oracle Database on premise or in the cloud, whether you prefer standard hardware or Oracle Engineered Systems, Oracle Cloud Infrastructure or Exadata Cloud Service.

If you are running your SAP on Oracle deployment on premise and would like to learn more about the benefits of moving SAP on Oracle Database to the Oracle Cloud, make sure to download and read the companion publication Oracle for SAP Cloud Update, which is available at <http://www.oracle.com/us/products/applications/ora4sap-cloud-update-4438361.pdf>

Sincerely,



Gerhard Kuppler

Oracle Vice President SAP Alliances



Oracle Database In-Memory

ORACLE DATABASE IN-MEMORY: BENEFITS



Innovation for Efficiency and Flexibility



Platform Continuity



Easy to Understand and to Deploy

INNOVATION FOR EFFICIENCY AND FLEXIBILITY




Make Existing Queries Faster :
Process reports and queries with sub-second response



Implement Simplified: SAP Data Models
e.g. Flat Cubes



Ask New Questions:
Queries that took too long in the past



Simplify Customizations:
e.g. get rid of user-defined indexes

PLATFORM CONTINUITY



No Infrastructure Changes:
Runs on all cloud and on-premise platforms



No Data Migrations or Reorganizations:
Storage location and format unchanged



No Feature Conflicts:
Can be combined with compression, encryption, etc.



No Lack of People and/or Knowledge:
Existing know-how still valid

ORACLE DATABASE IN-MEMORY

The Concept

Oracle Database 12c comes with a Database In-Memory option, however it is not an in-memory database. Supporters of the in-memory database approach believe that a database should not be stored on disk, but (completely) in memory, and that all data should be stored in columnar format. It is easy to see that for several reasons (among them data persistency and data manipulation via OLTP applications) a pure in-memory database in this sense is not possible. Therefore, components and features not compatible with the original concept have silently been added to in-memory databases such as HANA. Oracle has chosen the opposite strategy: Data can be populated into an In-Memory Column Store whenever this makes sense. In all other cases, data is stored and handled as it has always been.

Innovation for Efficiency and Flexibility

Oracle Database In-Memory adds in-memory functionality to Oracle Database for transparently accelerating analytic queries by orders of magnitude, enabling real-time business decisions.

Oracle Database In-Memory accelerates both Data Warehouses and mixed workload OLTP databases and is easily deployed under any existing application that is compatible with Oracle Database. No application changes are required. Database In-Memory uses Oracle's mature scale-up, scale-out, and storage-tiering technologies to cost effectively run any size workload.

Platform Continuity

The new column format is a pure in-memory format. Tables are stored on disk using Oracle's existing row-based formats. Since tables as such are never stored in columnar format on disk, there are no additional storage costs or storage synchronization issues.

Nor is there a need to modify the database. Oracle Database 12c In-Memory can be implemented without a database migration or a table reorganization.

As a result, the new Oracle Database 12c In-Memory feature is fully compatible with existing standard or optional database features such as table and index compression, table encryption, and table partitioning. It is also compatible with the scale-out architecture provided by Real Application Clusters (RAC) and with all existing high availability technologies (such as Data Guard).

Easy to Understand and to Deploy

In addition to being compatible at the database feature and 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.

In addition, however, Oracle customers expect mechanisms, which allow for fine-grained control and tuning. Oracle Database 12c In-Memory provides such mechanisms. E.g.: Tables can contain "cold" data, which are neither updated anymore nor accessed by queries. If those tables are very large, it would be a waste of memory to keep them completely in the in-memory column store.

NEW IN
18^c

IN-MEMORY SCAN OPTIMIZATIONS

In-Memory Optimized Arithmetic: Number columns are encoded using an optimized (binary) format that allows aggregations and other arithmetic operations to be performed in hardware. Scan performance may benefit by up to a factor of 9.

In-Memory Dynamic Scans further increase scan performance. They use idle CPU resources automatically to scan data in parallel and maximize CPU usage. Dynamic scans use multiple lightweight threads of execution within a process, and this helps keep the performance overhead low.

NEW IN
12.2

IN-MEMORY FAST START

Ordinarily, when an instance is restarted, the in-memory column store must be rebuilt from scratch, a process referred to as in-memory populate. This process can be CPU-intensive, since it must convert row-format data into compressed columnar data. With Oracle Database 12.2, the In-Memory Fast Start mechanism can significantly reduce the total time required for population by keeping a checkpointed copy of the column store on disk. As a result, when the instance is restarted, the checkpointed copy can be read back into memory without requiring any transformation of the data.

SAP Application Optimization Support

APPLICATION OPTIMIZATION TYPES



Workload Distribution Optimization



Data Model Optimization



Application Development Optimization

WORKLOAD DISTRIBUTION OPTIMIZATIONS



ABAP Core Data Services (CDS):


Push data-intensive computations from application layer to database layer

DATA MODEL OPTIMIZATIONS



Flat Cubes:

Optimize SAP BW Cubes for In-Memory Computing (SAP Note 2335159)



Declustering/Depooling:

Convert clustered and pooled tables to transparent tables (SAP Notes 1892354 and 1835008)

APPLICATION DEVELOPMENT OPTIMIZATIONS



CDS-based Application Development:

Build your own Fiori applications based on CDS views (Whitepaper "ABAP Core Data Services on any DB")

SAP APPLICATION OPTIMIZATION SUPPORT

Workload Distribution Optimizations

SAP used to think of a database as a dumb data store. Whenever a user wants to do something useful with the data, it must be transferred, because the intelligence sits in the SAP Application Server.

The disadvantages of this approach are obvious: If the sum of 1 million values needs to be calculated and if those values represent money in different currencies, 1 million individual values are transferred from the database server to the application server – only to be thrown away after the calculation has been done.

As a response to this insight, SAP developed the „Push down“ strategy: push down code that requires data-intensive computations from the application layer to the database layer. They developed a completely new programming model that allows ABAP code to (implicitly or explicitly) call procedures stored in the database. And they defined a library of standard procedures, called *SAP NetWeaver Core Data Services (CDS)*.

20 years earlier, Oracle had already had the same idea and made the same decision. Since version 7 Oracle Database allows developers to create procedures and functions that can be stored and run within the database. It was therefore possible to make CDS available for Oracle Database as well, and today SAP application developers can make use of it.

Application Development Optimizations

The benefits of the CDS framework described above are by no means restricted to SAP applications (i.e. standard applications created by SAP developers). For customers, home-grown applications are an essential part of their SAP landscape. Many of these applications could significantly benefit from using CDS features.

CDS views can be exposed via OData. Based on the OData exposure of CDS, it is then rather straightforward to create *SAP Fiori applications* using the development framework SAP WEB IDE.

Data Model Optimizations

A third important aspect is the optimization of data models. In many cases SAP's data models had been defined 15 or 20 years ago and optimized for disk-oriented databases. But, as it turned out, what was an optimization in the age of disk-based computing is an obstacle in the age of in-memory computing.

The most famous example is probably the internal structure of an SAP BW cube. From a business or user perspective what looks like one single “cube”, is actually a set of multiple tables, and the relationships between them can be described as a multi-level hierarchy. But this complex structure, which requires many joins when a query or a report is executed, slows down in-memory databases considerably. Therefore, SAP designed a new, simpler data model for SAP BW on HANA and consequently called it *HANA-Optimized InfoCubes*. However, this new data model is not only optimized for HANA. It is optimized for in-memory computing in general. Therefore SAP on Oracle users who have activated Oracle Database In-Memory can implement it as well, the only difference being the name (*Flat InfoCubes* or simply *Flat Cubes*).

A less famous, yet important optimization is Table Declustering. A cluster table stores a complete (logical) record in one single (physical) table column. Such a complex value can be interpreted by the SAP Application Server, but not by a database server – which means that code pushdown is not possible, if a cluster table is involved. Therefore SAP now supports *Table Declustering*, for HANA as well as for the Oracle Database.

NEW

SUPPORT FOR SEMANTICALLY PARTITIONED OBJECTS (SPOs)

SPOs are InfoProviders that consist of several InfoCubes or DataStore objects with the same structure. They can be said to be divided into partitions on the SAP level (as opposed to table partitioning on the Oracle Database level). Support for SPOs was not included in the first implementation of Flat Cubes, but added in November 2018.

Deployment Flexibility

DATABASE FEATURES AND OPTIONS



HW/OS Platform Independence




DB Architecture Options



Cross-Platform Administration

HARDWARE AND OPERATING SYSTEM PLATFORM SUPPORT



Generic Servers:
Unix, Linux, Windows



Oracle Cloud:
Oracle Cloud Infrastructure,
Exadata Cloud Service



Oracle Engineered Systems:
Exadata, SuperCluster,
Database Machine,
Exalogic



Virtual Machines:
VMs on generic servers and
Oracle Engineered Systems;
Virtual Shapes in the Oracle
Cloud

ORACLE DATABASE ARCHITECTURES



Single-Instance:
One instance,
one database



Real Application Clusters:
Multiple instances
on different nodes,
one database



Multitenant:
One instance,
multiple databases



Data Guard:
Standby Database for HA
(can be combined with single-
instance, RAC, Multitenant)

DEPLOYMENT FLEXIBILITY

Hardware and Operating System Platform Options

Generic Servers: Oracle Database is available for all major operating systems and can run on server hardware provided by many different vendors.

Oracle Engineered Systems: While each of the IT infrastructure layers provides leading-edge technology in itself, Oracle went one step further and designed engineered systems that are pre-integrated to reduce the cost and complexity of IT infrastructures. *Oracle Exadata Database Machine, Oracle Database Appliance, Oracle SuperCluster and Private Cloud Appliance* are designed to achieve enterprise performance levels that are unmatched in the industry.

Oracle Cloud: Customers who want to run their SAP on Oracle systems in the Oracle Cloud, can choose between two different services:

- *Oracle Cloud Infrastructure (OCI)* provides compute and storage offerings to run most demanding workloads in a secure and highly available cloud environment.
- With *Exadata Cloud Service*, customers can run Oracle databases in the cloud on Exadata, with the same extreme performance and availability experienced by thousands of organizations deploying Exadata on premise.

Virtualization is an implementation option that can be combined with all platform alternatives just discussed. On generic servers and Oracle Engineered Systems, Oracle Database Server instances can run in virtual machines. Oracle Cloud supports Virtual Shapes.

Oracle Database Architecture Options

Single Instance: Consisting of one single Oracle database and one single instance, this architecture is easy to configure and manage. However, its scalability and high availability options are limited.

Real Application Clusters (RAC) combines workload distribution, scalability, high availability, better manageability, and cost savings. This architecture allows multiple instances to access the same database at the same time. As these instances run on different

machines, customers have the option to implement a scale-out approach: 4, 6, or 8 small servers can handle the workload instead of one big and more expensive server. If one of the RAC servers fails, one of the other instances can take over.

Oracle Multitenant is an Oracle Database 12c option that helps customers reduce IT costs by simplifying consolidation, provisioning, upgrades, and more. It relies on a new architecture that allows one single instance to manage multiple pluggable databases (PDBs), which are consolidated in a container database (CDB). Multiple existing, independent databases may be converted to PDBs and consolidated into a single CDB. From an application's point of view, nothing has changed, as the PDB is still its database. This means: No application changes are required to adopt this architecture.

Data Guard: Data Guard can be combined with all of them. RAC provides high availability by multiplying the number of Oracle instances. However, there is still only one database. Data Guard removes this single point of failure. The technology allows customers to set up a standby (shadow) database as a copy of the primary (production) database and then keep the two databases synchronized.

NEW IN
18c

POLYMORPHIC TABLE FUNCTIONS

Because of its feature richness, Oracle Database frequently provides customers a choice between different ways to execute an SAP application request. Polymorphic Table Functions are a new option for optimizing the execution SAP's FOR ALL ENTRIES clause.

NEW IN
12.2


OPERATING SYSTEM SUPPORT

Oracle Database 12c Release 2 (12.2) is the first release certified on Microsoft Windows Server 2016.

Oracle Database 12c Release 2 (12.2) certified on SUSE Linux Enterprise Server (SLES) 15.

Very Large Database Support

DATABASE FEATURES AND OPTIONS



ORACLE ADVANCED COMPRESSION: BENEFITS



Table and Index Compression:
(1) Reduce disk space needed for database by 50% or more



Table and Index Compression:
(2) Improve performance by reducing I/O and storing more data in memory



Information Lifecycle Management (ILM):
Implement automatic deferred table/partition compression

TABLE AND INDEX PARTITIONING: BENEFITS



Improve Application Performance:
Reduce runtime of batch jobs (e.g. monthly reports)



Improve Database Manageability:
(1) Avoid fragmentation as a consequence of SAP Archiving



Improve Database Manageability:
(2) Foundation for additional ILM options

VERY LARGE DATABASE SUPPORT

Advanced Compression

Oracle Advanced Compression (as introduced in *Oracle Database 11c*) uses a different format for storing table and index data. Use of this compressed format helps reduce the database size by 50% or more, which is the *essential benefit* in the sense that it is the effect Advanced Compression has been designed for. A smaller source database footprint also means that the creation of back-ups and other copies requires less time.

As an *additional benefit* customers using Advanced Compression may see a performance improvement. Additional (as opposed to essential) here means: It may, but it is not guaranteed to happen.

Oracle Database 12c Advanced Compression adds two new features: Heat Map automatically tracks modification and query timestamps, providing detailed insight into how data is being accessed. *Automatic Data Optimization (ADO)* automatically moves and compresses data according to user-defined policies based on the information collected by Heat Map.

Heat Map and Automatic Data Optimization allow you to introduce a new parameter: If a table or partition should be compressed, when would you like it to be compressed? In Oracle Database 11g compression happens immediately or not at all. In Oracle Database 12c you can specify that data should be loaded today and (automatically) compressed a day or a month later.

Hybrid Columnar Compression (HCC)

In addition to algorithms for compression of structured and unstructured data, Oracle Exadata and Oracle SuperCluster support Hybrid Columnar Compression. This technology utilizes a combination of both row and columnar methods for storing data. The hybrid approach achieves the compression benefits of columnar storage, while avoiding the performance shortfalls of a pure columnar format. The compression ratios that can be achieved by using HCC are much higher than those seen with “normal” compression.

Table and Index Partitioning

In more and more situations today the distribution of the data on disk turns out to be a problem:

- Single queries or complex batch jobs accessing a certain subset of the table data need too much time to complete.
- Data load (SAP BW) is either slow, because it must update many indexes; or indexes are dropped and

rebuilt, in order to reduce load time, but this slows down user queries.

- Data archiving results in heavily fragmented databases.
- Customers want to implement information lifecycle management.

Oracle Partitioning divides tables and indexes into smaller units (called partitions) and forces all data to be stored in the appropriate unit. Partitions can be accessed and managed individually and independently from each other. Therefore:

- Ideally a query now finds all relevant data in one single partition and can ignore all other partitions („partition pruning“). This can reduce the runtime considerably.
- If the indexes defined on a partitioned table are partitioned as well, individual index partitions can be dropped and rebuilt while all other partitions remain untouched.
- The data archiving strategy can be based on the partition structure, and this can avoid disk space fragmentation.
- Partitioning is one of the basis technologies for information lifecycle management.

Oracle Partitioning is certified for all SAP NetWeaver applications. It is configured and used by default in SAP BW on Oracle. In SAP OLTP systems, it can either be implemented using the SAP Partitioning Engine (which covers the data archiving issue) or by Oracle ACS for SAP.

NEW IN
12.2

ADVANCED INDEX COMPRESSION HIGH

Advanced Index Compression LOW, introduced with Oracle Database 12c Release 1 (12.1), aimed to automate index key compression. It automatically decides which indexes to compress and computes the prefix column count within compressed indexes.

New with Oracle Database 12c Release 2 (12.2), Oracle has introduced **Advanced Index Compression HIGH**, geared towards dramatically improving index compression ratios. It introduces many additional compression techniques, which improve the compression ratios significantly while still providing efficient OLTP access.

Database Security

DATABASE FEATURES AND OPTIONS



Oracle Database Security Features



Oracle Advanced Security



Oracle Database Vault

ENCRYPTING DATA IN TRANSIT AND AT REST



Network Encryption:

Encrypt data travelling between SAP Application Server and Oracle Database Server (Oracle Database)



Transparent Data Encryption:

Encrypt data in production database files (Advanced Security)



Backup Encryption:

Encrypt data in backup copies (Advanced Security)

ORACLE DATABASE VAULT: RESTRICTING ACCESS TO DATA



Highly Efficient and Flexible Privilege Management:

Use factors such as date, time, or IP address



Segregation of Duties:

Have more than one person required to complete a task



Access Policy for SAP Applications:

No need to start from scratch. Use predefined policy. Modify, if needed.

DATABASE SECURITY

In order to read or update data in an SAP on Oracle system, the obvious and only tool for legitimate users is the SAP application. Attackers, however, who want to bypass SAP's user management and access control, could use other tools in order to access data at different stages of the workflow. In those cases Oracle can help to protect the data.

Protecting Data in Transit

Attackers could use a network sniffing tool to capture data in transit or even to intercept and modify them. To solve this issue, the Oracle Database offers *Network Encryption*, which allows customers to encrypt data travelling between SAP Application Server and Oracle Database Server. Encryption itself protects data against unauthorized read operations, while crypto-checksumming ensures data integrity.

Protecting Data at Rest

There are two different places where data can be „at rest“. One of them is the production database, the other its backup copies. Both places can be protected using features of the *Oracle Advanced Security* option.

Oracle *Transparent Data Encryption (TDE)* is applied to data in the files which make up the production database. As the name indicates, TDE is transparent to the application; no application changes are required.

Transparent Data Encryption comes in two flavors. The older one (available since Oracle Database 10g) is called *Column Encryption*, because you select just a few of the many SAP tables, or even individual columns of these tables that contain particularly sensitive data and you encrypt them. Everything else remains unencrypted. The newer one (available since Oracle Database 11g) is called *Tablespace Encryption*. It allows you to encrypt complete tablespaces, which may contain hundreds, thousands, or tens of thousands of tables. For Oracle Database 12c or later it is recommended to use tablespace encryption only. Starting with SAP NetWeaver version 7.20, BRSPACE can be used to set the encryption attributes on the tablespace level.

However, it is generally much easier to steal *backups of the database* files than the production database files themselves. If you simply backup your database files, only those tables (or columns) that are encrypted in the production database files are encrypted in the backup copies. *Backup Set Encryption* increases the protection level by encrypting the whole backup copy.

Protecting Data against Unauthorized Access of Privileged Users

Data encryption does not help, if attackers do not use third-party tools, but Oracle tools to bypass the SAP applications and connect directly to the database. This is particularly dangerous in the case of *privileged database users* (database administrators). And it is particularly relevant, if database administration is outsourced or data are stored in the cloud.

Oracle Database Vault replaces the traditional database privilege management strategy with a new, more flexible and more powerful one. It goes far beyond traditional user–privilege or user–role correlations. Oracle Database Vault allows companies to implement and enforce concepts such as the segregation of duties or the four eyes principle.

Oracle Database Vault for SAP is not just a toolbox. It comes with a predefined security policy that can be used as is or modified.

Auditing

The implementation of data encryption and access policies is only one side of data protection. It must be supplemented by monitoring of the user actions. Oracle Database provides a depth of auditing that readily enables system administrators to implement enhanced protections, early detection of suspicious activities, and finely-tuned security responses.

Oracle Database 12c *Unified Auditing* enables selective and effective auditing inside the Oracle database using policies and conditions. The new policy-based syntax simplifies management of auditing within the database and provides the ability to accelerate auditing based on conditions.

NEW IN
12.2

ONLINE TABLESPACE ENCRYPTION

In previous releases, only new tablespaces could be encrypted. Existing data had to be exported and re-imported. Oracle Database 12c Release 2 allows encryption of existing tablespaces while they are online and in readwrite mode.

Starting with Oracle Database 12.2, it is also supported to encrypt Oracle-supplied tablespaces (SYSTEM, SYSAUX, etc.) in addition to the tablespaces containing user/application data.

Manageability

MANAGEABILITY CONCEPTS



Integrated Management Tools



**Oracle Multitenant:
Manage many databases
as one**



**Autonomous Database
Strategy**

ORACLE MULTITENANT: MANAGE MANY DATABASES AS ONE



Better Resource Utilization and Resource Management:

General processes and data implemented only once in Container Database



Simplified Patching:

Patching is done in Container Database layer



Simplified Cloning and Provisioning:

Container Database (CDB) as infrastructure for administration of Pluggable Databases (PDBs)

INTEGRATED MANAGEMENT TOOLS: EXAMPLES



SAP BR*Tools:

Interactive interface for Oracle Database management



SAP DBA Cockpit:

Platform-independent tool for database monitoring and management



SAP Repartitioning Tool:

Tool for SAP BW table repartitioning and migration to Flat Cubes

MANAGEABILITY

Integrated Administration Tools

From the very beginning, SAP has not only provided interfaces for end users, but also tools for administrators responsible for managing SAP-related databases. The intention has always been to hide the complexity which necessarily arises if an application with many strict requirements meets a database that comes with many features and numerous options.

Examples are:

- **Software Provisioning Manager (SWPM)** is mainly an installation tool. However, it is also used to activate and configure optional Oracle Database features.
- **BR*Tools** is a set of tools for Oracle Database management in SAP environments (disk space management, memory management, backup, recovery, etc.).
- **DBA Cockpit** is a platform-independent tool mainly designed for monitoring and control of databases in SAP environments.
- **Repartitioning Tool**, as its name suggests, had originally been designed for repartitioning of partitioned tables. Meanwhile it is also used for tasks such as conversion of traditional SAP BW cubes to Flat Cubes.

Oracle Multitenant – Manage Many Databases as One

Many SAP landscapes consist of a few large and a considerable number of small or very small systems. In addition, almost all SAP landscapes require multiple identical copies of certain databases (development and test systems). However, the existence of many small and/or multiple identical database systems running on as many independent database servers has several disadvantages:

- Many small systems (even virtualized ones) use too many hardware resources (memory, CPU).
- Too much time is spent for the administration of so many small database systems.

Oracle Multitenant reduces resource consumption by separating one single “Container Database” (CDB) and multiple “Pluggable Databases” (PDBs). It simplifies administration by moving standard operations to the “container database” level. This new architecture delivers all the benefits of managing many databases as one, yet retains the isolation and resource prioritization of separate databases.

Sharing of background processes, memory structures, system-wide metadata, and database files results in considerably decreased resource consumption. In addition, Oracle Database 12c Resource Manager is extended with specific functionality to control the competition for resources between PDBs within a CDB.

By consolidating existing databases as pluggable databases, administrators can manage many databases as one. E.g.:

- Instead of upgrading or patching every single database (PDB), it is now sufficient to upgrade or patch the CDB. All hosted PDBs are upgraded „in-place“.
- Instead of executing separate database backups, administrators only require to back up their database at the CDB level. All PDBs consolidated into that container will be backed up as one. Nevertheless, administrators retain the flexibility to perform recovery operations at individual PDB level, if required.
- Administrators maintaining standby systems in another data center (using Data Guard or Active Data Guard) will only need to set up a single standby configuration at the CDB level. This allows all PDBs consolidated in that container to be replicated.

Autonomous Database Strategy

Oracle Autonomous Database is a cloud-based technology designed to automate many of the routine tasks required to manage Oracle databases. The technology combines Oracle Database 18c with a set of automated administration services that use machine learning algorithms. The combination is offered as a cloud service called Oracle Autonomous Database Cloud.


The autonomous database is the logical consequence of Oracle’s research and development efforts during the last 10 years. Step by step features that had to be configured manually in the past are now managed automatically (automatic disk space management, automatic memory management, etc.).

It is planned to make this service available for SAP customers in the future.

Oracle Database for SAP

Same benefits ...

... in the Cloud:



Best Performance & Scalability



Best Deployment Flexibility



Best Availability & Reliability



Best Support for DB Consolidation



Best Support for Very Large Databases



Best Database Security



Best Manageability & Self-Management



Best Integration of Hardware & Software

... on Premise:



Best Performance & Scalability



Best Deployment Flexibility



Best Availability & Reliability



Best Support for DB Consolidation



Best Support for Very Large Databases



Best Database Security



Best Manageability & Self-Management



Best Integration of Hardware & Software

ORACLE®

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores,
CA 94065 U.S.A.

Christoph Kersten
christoph.kersten@oracle.com

Altrottstr. 31
69190 Walldorf, Germany
Tel. ++49 (0) 6227 356-285
Fax++49 (0) 6227-3847366
Internet: www.oracle.com/sap
<https://www.sap.com/community/topic/oracle.html>

Copyright © 2019, 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. 0615 SAP® is a registered trademark of SAP – SAP NetWeaver is a registered trademark of SAP – SAP HANA® is a registered trademark of SAP