

ORACLE DATABASE FOR SAP: LATEST DATABASE TECHNOLOGY AND SUPPORT FOR APPLICATION OPTIMIZATIONS

Strategy and Roadmap

From the very beginning, the *Oracle Database for SAP* or *SAP on Oracle Database* strategy had been based on two pillars. The first pillar is the integration of Oracle Database features with the SAP environment. The second pillar is the integration of SAP application features with the Oracle database.

The need to *integrate Oracle Database features with the SAP environment* has always been visible. It was particularly obvious, when Oracle released new database features for which the SAP architecture was not prepared. An example that many customers still remember is the project to integrate Real Application Clusters (RAC) into an SAP architecture based on the assumption that there can be many SAP Application Server instances, but only one Database Server instance. The certification of Oracle Multitenant was a similar architectural revolution and required no less effort than the RAC certification.

The need to *integrate SAP application features with the Oracle Database*, on the other hand, has only rarely been recognized. The classic SAP applications (such as R/3 and BW) have been developed on the Oracle Database. Later on, when SAP started to support IBM DB2 and Microsoft SQL Server, they put the least common denominator strategy in place, i.e. they used only those database features that were available in all supported databases. Not much stress, therefore, on the Oracle Database.

This has changed with the advent of SAP's own database (HANA). SAP realized very soon that they had to drop the least common denominator strategy and change their applications: As long as SAP applications treat HANA as a database similar to all other databases, it is very difficult to convince customers that there is a benefit in implementing HANA. Therefore, SAP has embarked on an application optimization project in order to allow SAP applications to make use of special HANA features.

"Special HANA features", however, does not mean "HANA-only features". There is nothing in HANA that cannot be done by the Oracle Database as well. Therefore, the need to integrate SAP application features with the Oracle Database has recently become more visible.

Today, both pillars supporting the SAP on Oracle Database strategy are clearly visible and important: Whenever Oracle releases a major new database feature, a development effort is needed to integrate it into the SAP architecture as well as the installation, administration and monitoring tools provided by SAP. Whenever SAP releases a new application optimization, a similar development effort is needed to integrate it with the Oracle Database technology.

Oracle recognizes the value that the tight integration between the Oracle database and the SAP application brings to our customers. Oracle's continuing commitment for both pillars is evident through the comprehensive set of database features provided and for the special HANA optimizations currently supported such as Core Data Services and Oracle Optimized Flat Cubes.

SAP Application Optimization Support

Oracle Database New Features and Options

Oracle Database Version Support: Status and Roadmap

An overview of the versions that are currently available or will be made available in the near future can be found in the Oracle Database Roadmap infographic on page 5. – For additional details see SAP Notes 1174136 and 2606828.

Oracle Database 11g

Standard maintenance for Oracle Database 11g ended on January 31, 2015.⁽¹⁾ The extended support phase, has started as of February 2015. For version 11.2.0.4 Oracle offered extended support until December 31, 2018 at no additional cost. Paid extended support is available until December 31, 2020. As a result, Oracle Database 11.2.0.4 is available to SAP customers until December 31, 2020. – For more information see SAP Note 2098258.

Oracle Database 12c (12.1)

Standard maintenance for Oracle Database 12.1 ended on July 31, 2018. The extended support phase started on this date. For version 12.1.0.2, Oracle is offering extended support until July 31, 2019 at no additional cost and paid extended support until July 31, 2021. As a result, Oracle Database 12.1 is available to SAP customers as a fully supported version until the end of July 2021. – For more information see SAP Note 2428722

Oracle Database 12c (12.2)

Standard maintenance for Oracle Database 12.2 (12.2.0.1) will end on November 20, 2020. No extended support is offered.

New Oracle Release Strategy

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 numbers, the release year is now used to designate a software version (18c, 19c, etc.). These annual software releases are being made available to SAP on Oracle customers as well.

Oracle Database 18c

Oracle Database 18c has been certified in March 2019. Standard maintenance will end in Q2 2021. No extended support is planned.

Oracle Database 19c

Oracle Database 19c is the first long-term support release based on the new release strategy. SAP certification is planned for December 2019. Standard maintenance will end on March 31, 2023. Extended support will end on March 31, 2026.



Oracle Database Versions: New Features Overview

The term “base certification” has been coined during the certification of Oracle Database 12.1 for SAP. This process was split into several phases, the first of them being a certification of the new database version without any new option (base certification). In this article the term is used for different Oracle Database versions. The “base certification” section contains a discussion of basic, but important features which are not mentioned in the following sections.

Oracle Database 12c (12.1)

Oracle Database 12c (12.1) came with two major new options and important enhancements to an existing option

which will be discussed later in this article and in other articles: Oracle Database In Memory (see pages 11-12), Oracle Multitenant (see pages 16-17), and support for Information Lifecycle Management provided by Oracle Advanced Compression (see pages 15-16). In addition, the following new features should be noted:

- **Advanced Index Compression** is a new form of index compression which is more efficient than standard Index Key Compression and requires less administration effort. Advanced Index Compression works well on all supported indexes, including those that are not good candidates with Index Key Compression.

¹All statements in this section are as of July 2019. Please keep in mind that these dates are subject to change at any time.

- **Advanced Network Compression** can be used to compress the data to be transmitted at the sending side and then uncompress it at the receiving side to reduce the network traffic. Advanced Network Compression allows transmission of large data in less time. It improves SQL query response time and saves bandwidth (see SAP Note 2138262).
- Data Guard – the functionality needed to set up stand by databases – is included in Oracle Database Enterprise Edition. Active Data Guard is an extra option. In Oracle Database 11g it offers additional features such as Automatic Block Repair and Fast Incremental Backup. **Active Data Guard Far Sync** the main new feature with Oracle Database 12c, allows customers to combine high performance (a characteristic of asynchronous data shipping) and zero data loss (a characteristic of synchronous data shipping) across large distance WANs. For details see the article “Implementing a Data Management Infrastructure for SAP with Oracle Database Options and Packs” (“Data Guard and Active Data Guard” section) on page 26.
- Oracle Recovery Manager (RMAN) provides a comprehensive foundation for efficiently backing up and recovering the Oracle Database. **Cross Platform Backup and Restore** enables you to transport data across platforms by using full and incremental backup sets. To perform cross-platform backups using backup sets, the destination database must be Oracle 12c or later. This newly added feature simplifies platform migration and minimizes read-only downtime on the source database.
- While RMAN remains the most popular tool to perform Oracle Database backups, another commonly used method for taking database backups is to create a storage snapshot of all files in the database. **Storage Snapshot Optimization** enables you to use third-party technologies to take a storage snapshot of your database without the need for putting the database in BACKUP mode.
- **ACFS**, which had not been supported on Exadata before, is now supported, if you are running Grid Infrastructure version 12.1.0.2 or later. In SAP environments this can be used for SAP shared file systems (/sapmnt,
- **High Availability Network File Storage (HANFS)** for Oracle Grid Infrastructure provides uninterrupted service of NFS V2/V3 exported paths by exposing NFS exports on Highly Available Virtual IPs (HAVIP) and using Oracle Clusterware agents to ensure that the HAVIPs and NFS exports are always online.
- The UNION and UNION ALL operators in SQL statements connect two or more branches (e.g. subqueries): `<branch1> UNION <branch2>`. Traditionally, in queries of this type, the branches execute one after another, meaning that only one branch is executed at a certain point in time, followed by the next branch, and so on in a serial manner. Oracle Database 12c introduces **Concurrent Execution of Union and Union All Branches**. The ability to parallelize branches leads to much faster statement execution. In particular, it improves SAP BW performance.
- **Online Move Partition:** Starting with Oracle Database 12c, the ALTER TABLE ... MOVE PARTITION operation functions as a non-blocking online DDL command, while DML operations continue to execute uninterrupted on the partition that is being moved. Additionally, global indexes are maintained when a partition is moved, so that a manual index rebuild is no longer required.
- **Move Datafile Online:** Prior to Oracle Database 12c, moving datafiles has been an offline task. Oracle Database 12c includes an enhancement to the ALTER DATABASE command to allow datafiles to be moved online.
- **Rebuild Index-Organized Tables:** Because index-organized tables are stored as B-tree indexes, you can encounter fragmentation as a consequence of inserts, updates and deletes. However, you can use the ALTER TABLE ... MOVE PARTITION statement to rebuild an index-organized table and reduce this fragmentation. For more information see SAP Notes 1856270 and 2087004.
- On Microsoft Windows, Oracle Database 12c supports the use of an Oracle Home User, which can be specified at installation time. The Oracle Home User is introduced to host Oracle Services for greater security using a low privileged non-administrator account. For more information see SAP Note 1915302.

Oracle Database 12c (12.2)

For SAP environments, Oracle Database 12.2 does not come with new options. However, it provides numerous enhancements of existing features or options that are relevant for performance, manageability, and platform support. Examples are:

- 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.
- 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 read-write 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.
- Operating System Support:** Oracle Database 12c Release 2 (12.2) is the first release certified on Microsoft Windows Server 2016 and SUSE Linux Enterprise Server (SLES) 15. This includes support for the database client, server, and Oracle Real Application Clusters.

Oracle Database 18c

Similar to Oracle Database 12.2, Oracle Database 18c does not offer completely new options. Among the new features considered for certification in SAP environments are:

- In-Memory Optimized Arithmetic:** The Oracle Database NUMBER data type has high fidelity and precision. However, NUMBER can incur a significant performance overhead for queries because arithmetic operations cannot be performed natively in hardware. The In-Memory optimized number format enables native calculations in hardware for segments compressed with the QUERY LOW compression option.
- A Polymorphic Table Function (PTF)** is a new type of table function whose return type is determined by the arguments passed into the PTF. Useful when SQL developers and database administrators want to provide generic extensions which work for arbitrary input tables or queries, it perfectly matches the ABAP SELECT clause FOR ALL ENTRIES.



Base Certification and Application Optimization

Theoretically speaking, implementation of Oracle support for SAP application optimizations is an ongoing project that runs completely independent from the certification of Oracle Database versions. However, in some cases certain version-specific features may be or are required.

A particularly interesting example is discussed in SAP Notes 1835008 and 1892354: Several application optimizations implemented by SAP can only be used, if some tables traditionally implemented as cluster tables are declustered. As the data in these cluster tables is normally stored in a compressed manner by SAP, customers find that the tables can grow considerably when they are converted to transparent tables.

Unfortunately some of the declustered tables have more than 255 columns. In those cases Oracle Database 11g Advanced Compression could not help to reduce their size, because in this version structured table data compression (OLTP compression) was not supported for tables with more than 255 columns.

In Oracle Database 12c Advanced Compression, the 255-columns limit is removed, and the table compression without this limit is available for SAP customers since the base certification. Therefore, with the Oracle Database 12c (and higher) Advanced Compression option, it is now possible to compress and manage the data residing in these very wide tables.



Core Data Services

Many people believe that SAP's decision to abandon the least common denominator strategy and to optimize their applications for HANA in mind are seen as a threat by Oracle. And it is certainly true that in the SAP world HANA is a competitor of the Oracle Database. However, in many cases SAP's new application optimizations are greeted with a sigh of relief by Oracle employees as well as by Oracle customers. Taking SAP Core Data Services (CDS) as an example, it is easy to explain why.

The main questions behind Core Data Services are: What is a database? What can it do? And what can it not do?

The traditional answer to these questions claims that a database is nothing but a dumb data store. It is a container that can permanently store data, but that's it. Whenever a customer wants to do something useful with the data, it must be transferred to the application server, because the intelligence sits in the application server.

Traditional SAP applications are based on this very concept. The disadvantages 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. The network traffic caused by this approach is suboptimal and suffers with poor performance.

More than 25 years ago the developers of the Oracle Database asked: Wouldn't it be nice, if this sum could be calculated on the database server side? Would this not improve the performance dramatically? They came up with a different answer to the question what a database is: A database is not only a data store, it can also store and execute procedures working with the data – pieces of code that originally were part of the application running on the application server, but are now moved to the database server. So the application is split into two tiers, one of them running on the application server, the other one on the database server, and therefore the database server is an *application tier*.

The Oracle developers did not only ask questions or come up with a new concept. They also built a new database version that was able to store and execute database procedures (Oracle 7, released 1992).

However, at that time the Oracle Database was the only database that could be used as an application tier. Stored procedures were not part of the least-common-denominator feature subset, and therefore SAP declined to use them.

20 years later, SAP started to promote HANA. One of the first things they discovered was that their own applications were the worst enemies of the new in-memory database architecture. If an application believes that a database is essentially a dumb data store, that only itself can do calculations efficiently and therefore individual values need to be transferred over the network, actively destroys all potential benefits of an in-memory database. At that time, SAP realized that they had to abandon the least common denominator strategy and its counterpart, the dumb data store concept.

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 in order to prevent pure chaos, they defined a library of standard procedures. This library is called Core Data Services (CDS). And they agreed to make this library available for non-HANA databases, too, if those databases support stored procedures.

The 20 years between the release of Oracle 7 and the release of SAP Core Data Services explain the sighs of relief breathed by Oracle customers and employees: The performance gains achieved by SAP’s push-down strategy would have been possible 20 years earlier. Better late than never.



Oracle Database In-Memory

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 are stored and handled as it always has been. (For more information on the concepts of Oracle Database In-Memory see the article “Implementing a Data Management Infrastructure for SAP

with Oracle database Options and Packs”, in particular the sections “Oracle Database In-Memory”, page 23 and “Summary”, page 35).

Oracle Database In-Memory has been certified for SAP in June 2015. Unlike similar options offered by competitors, the use of Oracle Database In-Memory is not limited to SAP Business Warehouse (SAP BW). It is supported for all SAP applications based on SAP NetWeaver, including typical OLTP applications. However, this does not mean that it is always a good idea to use Oracle Database In-Memory. This option is a solution for a specific problem – or for a certain set of problems. It cannot solve all problems. It cannot improve performance in all cases. If used in an inappropriate manner, it can even – like a pure in-memory database – degrade system performance. Therefore, the SAP applications that can benefit from data being loaded into the column store must be selected carefully.

Applications must be selected, individual tables must be selected – the implementation of Oracle Database In-Memory in SAP environments seems to be difficult. However, early adopters consistently mention as their very first experience that Oracle Database In-Memory for SAP can be implemented quickly and easily. This seems to be counterintuitive, but it is not.

First, many customers are already aware of the queries and jobs that take too much time to complete, and they know which tables are involved. In those cases the task to select appropriate SAP applications and tables is trivial.

Second, for customers who do not want to implement Oracle Database In-Memory in order to fix specific issues, but prefer a general approach, Oracle provides an In-Memory Advisor – a wizard that analyzes the workload of a particular system and recommends tables to be populated into the column store based on the amount of memory that is available. (This means that the frequently asked question “How

much memory do I need in order to use Oracle Database In-Memory?” is completely meaningless. It’s the other way around: You tell Oracle how much memory you have, and the advisor will let you know how that amount of memory can be used in the most efficient way.)

Third, once the relevant tables are determined, everything is easy and breathtakingly fast: By issuing an `ALTER TABLE XXX INMEMORY` statement you declare that those table data should be available in the column store and from this point on everything else happens automatically in the background.

Finally, unlike the migration to an in-memory database such as HANA, the implementation of Oracle Database In-Memory does not require a revolution: no new hardware, no new operating systems, no new database. Customers can continue to use the existing infrastructure, and what administrators need to know about Oracle Database In-Memory can be learned within a few hours.



Flat Cubes

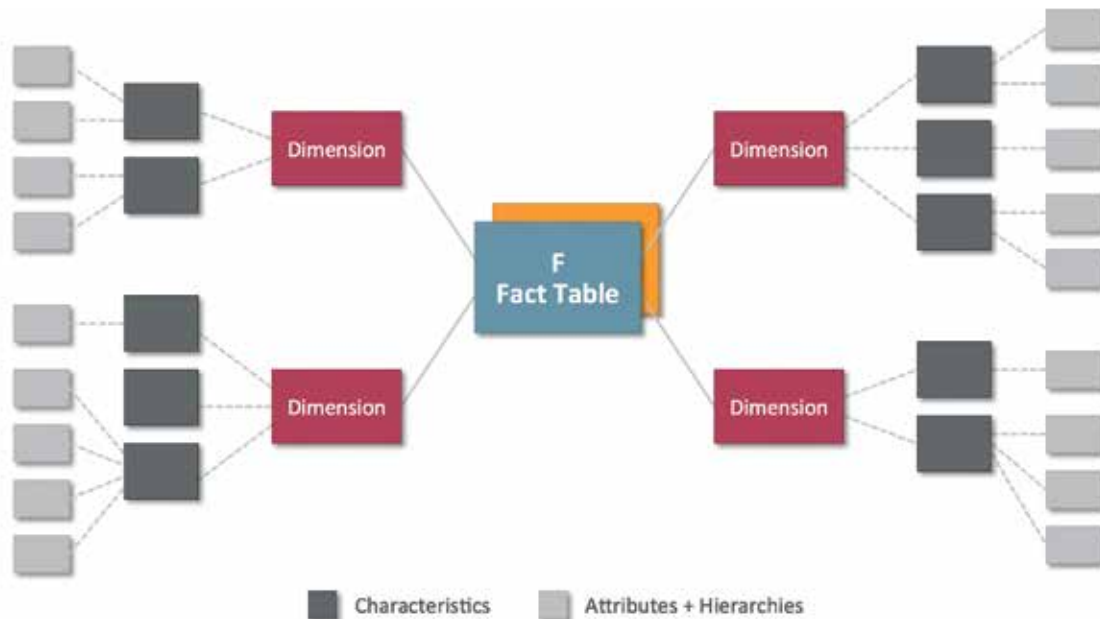
When the certification of Oracle Database In-Memory for SAP was announced in June 2015, the announcement included a couple of restrictions. In particular, it was strongly recommended not to drop any standard indexes or aggregates. This caused some disappointment, because from a pure Oracle perspective indexes are not needed anymore when the base tables are populated into the column store, and can therefore be dropped.

But in this case (as in all other cases described in this article) the Oracle/SAP development team, which is responsible for the integration of SAP and Oracle technologies, had to follow SAP’s learning curve. The situation immediately after the certification of Oracle Database In-Memory for SAP (in this case: for SAP BW) simply mirrors the early stages of SAP’s project to provide SAP BW on HANA.

The disappointment mirrors SAP’s experience that the traditional SAP BW data model is not compatible with the new concept of an in-memory database. And Flat Cubes,

which will be explained in this section, mirror the new data model that SAP designed for HANA.

In many cases, data to be loaded into the Business Warehouse arrive *as very wide records*. E. g. company name, zip code, city, and street address are combined with carrier details, order number, order date, invoice number and dozens, if not hundreds of other data items in one single record. But in the early days of data warehousing, when databases were disk-based only and disk space was expensive, it was not acceptable to waste disk space for redundant data such as the company or the carrier details which occur 1000 times, if that particular company sends 1000 items, and 100,000 times, if that particular carrier is engaged to fulfill 100,000 shipments. Therefore database architects came up with a design *called star schema*: subsets of data which belong together (all customer details, all carrier details) are moved to separate tables, which are called *dimension tables*. The remaining data plus IDs pointing to the relevant entries in the *dimension tables* is stored in the *fact table*.



Traditional "star" (= extended snowflake) schema

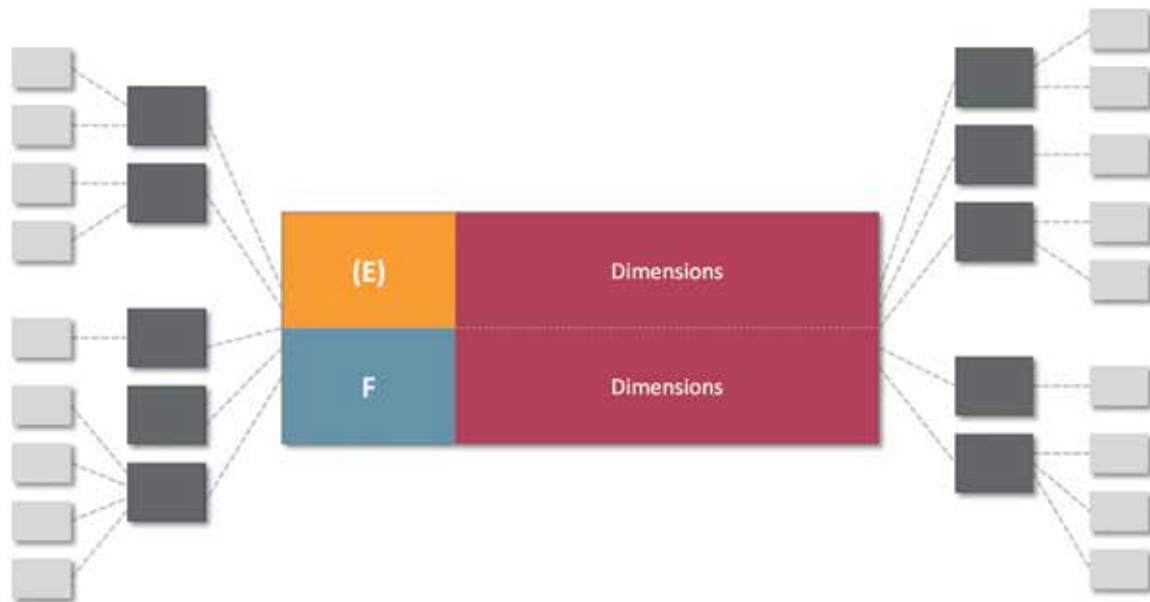
Such a split was not enough in all cases. E. g. a certain combination of zip code, city name and street may occur several times in the CUSTOMERS as well as in the CARRIERS table. If the same split operation is applied again, additional tables are created which, however, are not connected to the fact table, but to the dimension tables. This results in a more complex, but also (from a disk-space point of view) more efficient design, which is called *snowflake schema*. High-end data warehouses such as SAP BW add yet another level of detail tables, thus relying on the *extended snowflake schema*.

This complex architecture has been designed in order to optimize the data model for the requirements of traditional, disk-only relational databases. However the new databases with their focus on memory – and in this respect there is no difference between SAP HANA and Oracle Database in-Memory – have very different requirements.

Therefore, SAP designed a new data model for SAP BW on HANA and consequently called it *HANA-Optimized InfoCubes*.

The simplest, but somewhat surprising description of HANA-Optimized InfoCubes is this: If the process of optimizing the SAP BW data model for disk-oriented databases led from flat and therefore wide records to the extended star schema, the process of optimizing the data model for memory-oriented databases is simply the way back from extended star to flat and wide.

Back but not all the way. HANA-Optimized InfoCubes combine the fact table (actually: the E and F fact tables) and the dimension tables (first level of details) in one single table, whereas the small level 2 and 3 tables (characteristics, attributes and hierarchies) remain in place. This change is sufficient to improve performance and manageability considerably.



New flat cube design

This new data model removes the main disadvantages of the previous data model without sacrificing its benefits. It is no longer necessary to split the incoming, wide records in order to distribute them over many tables – this speeds up data load. The traditional indexes are not needed anymore – this speeds up data load as well. It is no longer necessary to join the tables later on – this speeds up query processing. And the main disadvantage of the flat data model, which originally motivated the development of the extended snowflake schema, the disk and memory requirements of redundant data, is no longer an issue thanks to the compression features that are available today for data on disk as well as for data in memory.

If this new data model is made available for a non-HANA database, “HANA-Optimized InfoCubes” is obviously not an appropriate name. “SAP BW Flat InfoCubes for Oracle” or simply “SAP BW Flat Cubes for Oracle” is exactly the same data model, called by a different name. It requires Oracle Database 12c or higher and Oracle Database In-Memory, as Flat Cubes outside of the Column Store do not make any sense.

Flat Cubes for SAP BW on the Oracle Database is generally available since June 2016. For more information see SAP Note 2335159.



Deferred Compression and Information Lifecycle Management

Some of the new features in Oracle Database 12c Advanced Compression have already been discussed in the “Base Certification Features” and “Base Certification and Application Optimization” sections. However, two major new features are still missing, because they were not included in the base certification, but certified a few months later (December 2015) for SAP environments: *Heat Map* and *Automatic Data Optimization (ADO)*. The basic concepts behind these two features are discussed in the article “Implementing a Data Management Infrastructure for SAP with Oracle database Options and Packs” (see in particular the section “Advanced Compression (Oracle Database 12c)”, page 21). Therefore, we will briefly look at the SAP-specific implementation details.

Oracle Database 12c Advanced Compression allows customers to distinguish between current (“hot”) and historical (“cold”) data. However, it is not clear what exactly the words “hot” and “cold” mean. So this needs to be defined:

```
ALTER TABLE <table_name> ILM ADD POLICY
<action>
AFTER <n> DAYS OF NO MODIFICATION;
```

The third line of this SQL statement answers the question. New data is considered “hot”. If it turns out that they have not been modified for a certain number of days (30, 60, 90 days), they are considered “cold” – assuming that the customer does not want to define intermediate levels such as “warm”. But if we look closer, we find that the only question that has been answered so far is: When do we call data “cold”? What we still do not know (and

what the database system still does not know) is: If data have cooled down – then what? What should happen? This is to be defined in line 2:

```
ALTER TABLE <table_name> ILM ADD POLICY
ROW STORE COMPRESS ADVANCED ROW
AFTER 40 DAYS OF NO MODIFICATION;
```

In this example we assume that (in this particular table) hot data is not compressed at all, and we tell the system that (a) any data not modified for 40 days should be considered cold and that (b) cold data should be compressed using the table compression algorithm provided by Oracle Database 12c Advanced Compression.

How do we, and how does the system know that data have not been modified for 40 days? It is the job of Heat Map to provide this kind of information. Heat Map automatically tracks modification and query timestamps at the row and segment levels, providing detailed insight into how data is being accessed. *Automatic Data Optimization (ADO)*, then, automatically moves and compresses data according to user-defined policies (such as that which we have used here as an example) based on the information collected by Heat Map.

So far the ALTER TABLE statement has been used to define the ILM policy. In SAP systems where we have to deal with tens of thousands of tables, this approach would be very cumbersome. Therefore the BR*Tools (BRSPACE) use a different option provided by the Oracle Database:

```
ALTER TABLESPACE TSX DEFAULT ILM ADD POLICY
ROW STORE COMPRESS ADVANCED ROW
AFTER 40 DAYS OF NO MODIFICATION;
```

In this example we do not define a special policy for an individual table, but a default policy on the tablespace level. It is automatically applied to all tables created in this tablespace, unless a table comes with an individual policy.

Customers running Oracle Database 12c or higher an Oracle Engineered System (Exadata, SuperCluster) can benefit from Hybrid Columnar Compression – a set of compression algorithms designed for purely historical data as an alternative to archiving. If Advanced Compression compresses data by a factor of 2 or 3, Hybrid Columnar Compression can easily achieve compression factors of 10 or 15.

In this situation, we would call data not modified for 40 days “warm”, and we would reserve the word “cold” for

data not changed during a considerably longer period (e.g. 6 or 12 months). We would keep the previous policy as compression tier 1 (form warm data) and add an additional policy as compression tier 2 (for cold data). And we would separate unpartitioned and partitioned tables in different tablespaces, because Hybrid Columnar Compression compresses complete partitions instead of individual blocks:

```
ALTER TABLESPACE TSY DEFAULT ILM ADD POLICY
ROW STORE COMPRESS ADVANCED ROW
AFTER 40 DAYS OF NO MODIFICATION;
```

```
ALTER TABLESPACE TSY DEFAULT ILM ADD POLICY
COLUMN STORE COMPRESS FOR QUERY LOW ROW
LEVEL LOCKING SEGMENT
AFTER 6 MONTHS OF NO MODIFICATION;
```



Oracle Multitenant

Oracle Multitenant is a new option for Oracle Database 12c that helps customers reduce IT costs by simplifying consolidation, provisioning, upgrades, and more. It is supported by a new architecture that allows a *container database (CDB)* to hold and manage many *pluggable databases (PDBs)* (see article “Implementing a Data Management Infrastructure for SAP with Oracle Database Options and Packs”, in particular section “Oracle Multitenant”, page 27).

With Oracle Multitenant, multiple existing databases may be converted to PDBs and consolidated into a single CDB. A PDB is a self-contained, fully functional Oracle Database. From an application’s point of view nothing has changed in any way, and that’s very important because it means that no application changes are required to adopt this architecture. From an application’s point of view, the PDB is the database. However, from an operational point of view the CDB is the database.

The CDB represents a single, consolidated operating environment. There is a single set of background processes and a single shared memory area (SGA), shared by all of the PDBs in the CDB. This architecture eliminates replication of overheads, making most efficient use of available resources. What this means is that you can minimize capital expenses (CapEx) because you can consolidate more applications per server (see figure 11a-e). From an operational perspective, you can manage all these consolidated PDBs collectively, greatly reducing operating expenses (OpEx). This applies to things like backups, configuration of high availability, application of patches and upgrades. These CapEx and OpEx reductions are part of how Multitenant delivers on the promise of Cloud computing.

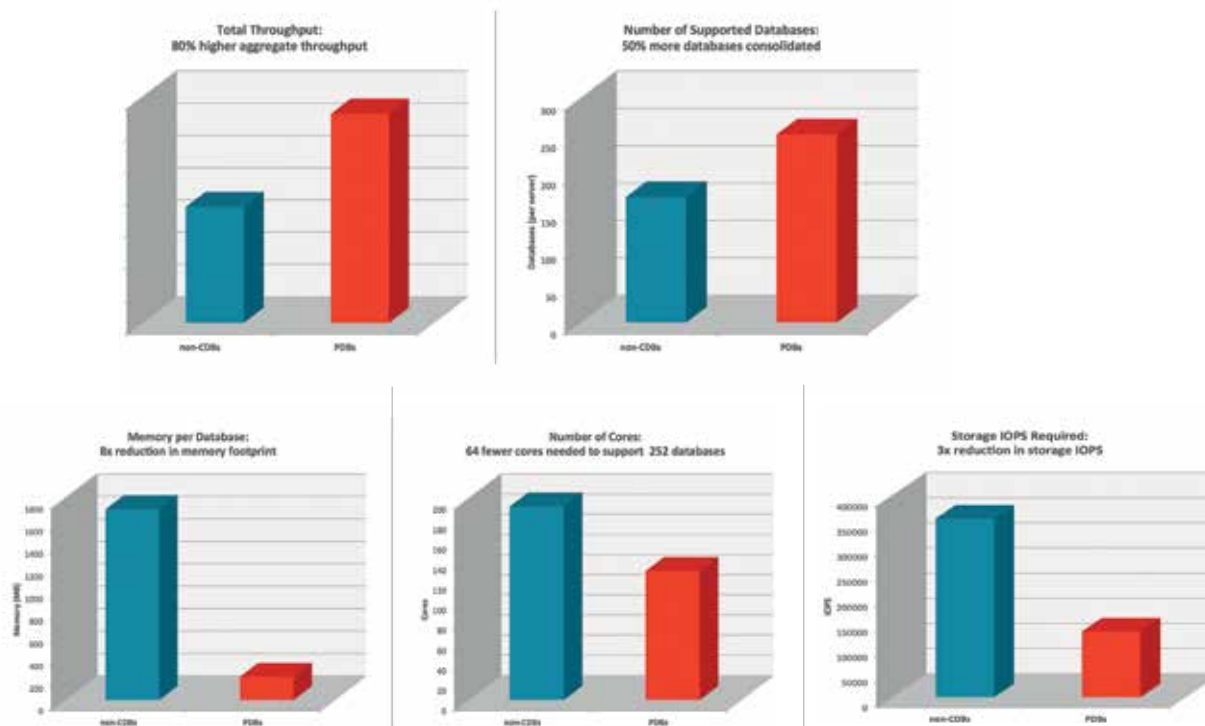
Oracle Multitenant is the architecture for the next-generation Database Cloud. Multitenant delivers true *economies of scale*. The expensive model of a VM containing a database is replaced by a pluggable database (PDB). Because there is negligible intrinsic cost to a PDB, the cost of each SAP system's PDB is reduced to the actual work they do.

The Oracle Multitenant architecture can be used for all SAP NetWeaver-based applications with the only exception

that a mix of SAP OLAP (BW) and SAP OLTP (ERP, CRM, ...) systems in the same container database is not supported.

For database administrators, the following tool support is available:

- As of version 1.0 SP 19, SWPM supports the creation of container databases (CDBs) and pluggable databases (PDBs). SWPM must be used for these tasks in order to guarantee compatibility of the created databases (directory paths, file names, etc.) with BR*Tools. For details see SAP Note 2336881.
- In most cases, customers will not create new databases, but convert existing stand-alone (non-CDB) databases to pluggable databases. SAP Note 2335850 describes the supported procedure for this kind of transformation.
- As of version 7.40 patch 24, BR*Tools support Oracle Multitenant. New configuration parameters, commands, and command options allow administrators to specify the target database(s) for operations initiated via the familiar BRCONNECT, BRSPACE, BRBACKUP/BRARCHIVE or BRRESTORE/BRRECOVER commands.



Consolidation performance test highlights: Oracle Multitenant achieves higher performance with less resource requirements