Oracle for SAP Database Update

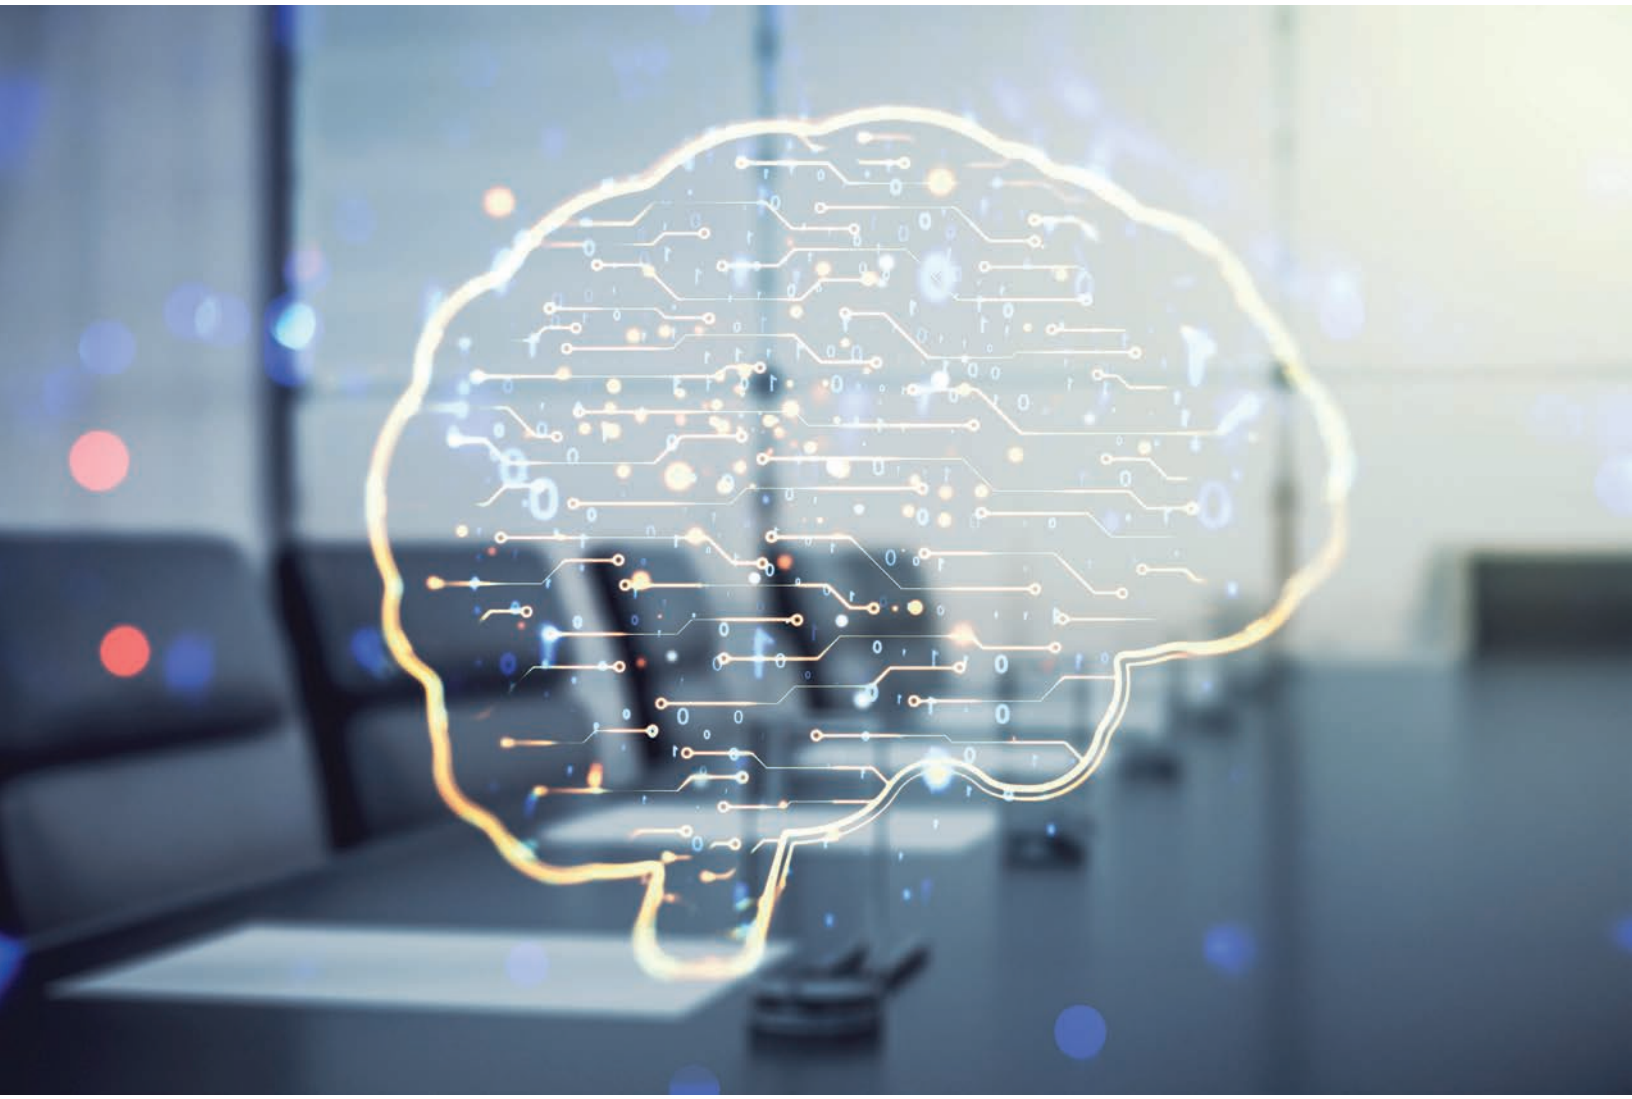

ORACLE

August 2024

Dear SAP Customer,

The relationship of Oracle Corporation and SAP SE has been and is based on a 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 30 years.

Our 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.

In 2024, Oracle released Oracle Database 23ai, the next long-term support release of the Oracle Database, with a focus on AI and developer productivity. Given the focus on AI in this release of the database, the database's name was changed from Oracle Database 23c to Oracle Database 23ai.

An overview of the versions that are currently certified or planned to be certified in the near future can be found page 15. For additional details read 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 on Oracle customers can enjoy these benefits. It does not matter whether one runs the Oracle Database on premise or in the cloud, nor whether one chooses standard hardware, Oracle Engineered Systems, Oracle Cloud Infrastructure, Exadata Cloud Services, or Exadata Cloud@Customer.

.

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 onto Oracle Cloud Infrastructure, make sure to download and read the companion publication "Oracle for SAP Cloud and Infrastructure Update", which is available at: http://www.oracle.com/sap

Sincerely,

Gerhard Kuppler

Oracle Vice President SAP Alliances

Table of Contents

Oracle Database for SAP News

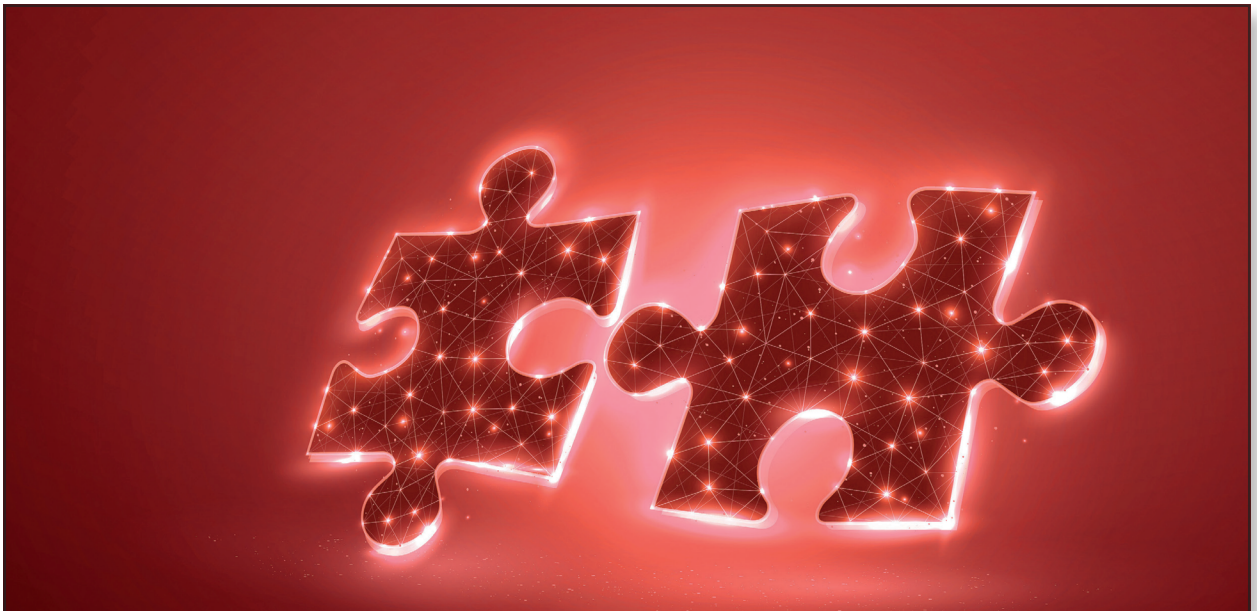

Interview with Gerhard Kuppler, Vice President SAP Alliances, Oracle

Working Closely to Benefit Thousands of Joint Customers

E3 Magazine spoke to Gerhard Kuppler, Vice President SAP Alliances at cloud provider Oracle, about topics such as collaboration with SAP, the Oracle database, and much more.

E3: How has the ECC/NetWeaver/Database relationship between Oracle and SAP evolved over the past five years?
Gerhard Kuppler, Vice President SAP Alliances, Oracle: Our collaboration dates back more than 30 years and continues to be strong, as evidenced by SAP's recently-announced commitment to extend support for SAP NetWeaver-based solutions until at least 2030. Further extensions cannot be ruled out given the large number of SAP customers still running SAP ECC and associated NetWeaver applications.

E3: There seem to be fewer activities between Oracle and SAP lately. Is there still an active collaboration between the two companies?
Kuppler: We continue to work closely to benefit our thousands of joint customers. We have a significant number of Oracle Database and OCI (Oracle Cloud Infrastructure) developers based in the SAP Development Center at Walldorf, as well as Oracle support engineers embedded in SAP teams around the world. We are collaborating closely to certify the latest versions of Oracle Database, Oracle Cloud, Oracle Engineered Systems, and Oracle Linux technology as these become available. For example, currently we are working to certify Oracle Database 23ai, the latest long-term release of Oracle's converged database.

E3: How do you perceive the global SAP community's development?
Kuppler: The results of the latest 2024 DSAG Investment Survey show that a large number of SAP customers are still running SAP ECC/NetWeaver. Many of these customers might have highly customized ECC systems which cannot be easily migrated to other ERP solutions, or they might not be able to justify the spend for a complete ERP system replacement. We will continue to support Oracle Database for these customers as long as SAP continues to support NetWeaver.

E3: Many SAP customers are beginning to move their systems to cloud providers, hyperscalers, and SAP itself. What would you recommend to SAP customers running ERP/ECC 6.0 with Oracle Database? How is the market developing?
Kuppler: Many forward-looking customers are choosing to move to the cloud. Oracle offers Exadata as a cloud service on OCI, so that customers running Exadata on-premises can easily move to Exadata Database Service on OCI.

E3: Are there other options?
Kuppler: For customers who need to maintain continuity or have specific requirements around data residency or latency-sensitive workloads, Oracle Exadata Cloud@Customer—where the entire Exadata cloud instrastructure is deployed in the customer's own data center—is a good choice. With this type of deployment, as with any cloud deployment model, the customer subscribes to the infrastructure without having to purchase the hardware.

Gerhard Kuppler,
Vice President SAP Alliances,
Oracle

E3: Can you give us a more detailed roadmap and some recommendations for SAP customers using Oracle Database?
Kuppler: SAP customers can look forward to the certification of Oracle Database 23ai to enable them to smoothly transition from Oracle Database 19c. Oracle Database 19c is currently under premier support until at least April 2026, and will be under extended support until at least April 2027. Therefore, customers would have sufficient time to upgrade their databases to the new long-term release.

E3: Where are you seeing more adoption of Application Continuity for Oracle Database?
Kuppler: We are seeing adoption by SAP customers who need to ensure zero disruption to active online users while performing downtime maintenance for one or more database nodes in RAC (Real Application Clusters) configurations.

E3: Will Oracle Database skills be still useful given SAP's move to RISE and Public Cloud editions of S/4 Hana?
Kuppler: Yes. Besides SAP NetWeaver applications, Oracle Database is used in thousands of non-SAP applications. Oracle Database skills are also portable between on-premises and cloud deployments of Oracle Database, and across different operating systems, hardware platforms, and hyperscaler environments beyond OCI.

E3: Could you specify which hyperscalers?
Kuppler: The sharp rise in adoption of Oracle Autonomous Database, the significant investments by Microsoft and Oracle in the Oracle Database@Azure service, and the recently announced investments by Google and Oracle in the Oracle Database@Google Cloud service, present opportunities for those with Oracle Database skills. Due to the complexity of RISE and of moving to Hana, many customers are considering alternative solutions.

E3: What would be your advice to customers still running ECC 6.0 today who have not yet decided on their next step?
Kuppler: Thousands of large enterprise customers still operate on ECC 6.0, and the application won't stop running come 2030. Our advice to these customers is to first stabilize and optimize their NetWeaver applications on a modern platform and long-term supported Oracle Database—whether Exadata on-premises, on OCI, or on Exadata Cloud@Customer. Once this is running smoothly, the customer can then assess and decide on next steps—whether that involves S/4 Hana or a different platform—without concerns for performance or stability issues in their ECC 6.0 environment.

E3: What do you make of the recent DSAG Investment Survey results?
Kuppler: SAP customers are a highly sophisticated audience. They demand to understand the concrete business benefits of migration before making a decision. I don't think it's helpful to confront them with an end-of-life deadline. Ultimately it is a business decision.

E3: What are the most significant benefits you would highlight to customers considering moving to the Oracle Cloud Infrastructure and why?
Kuppler: The biggest benefit is peace of mind. We currently run the world's largest ECC IaaS deployment on OCI. There is no need to worry about performance, as we have Exadata technology available in the cloud. We also offer reliability as we offer the only SAP-supported RAC setup in the cloud, as well as value with the inclusion of the Oracle Linux Premier Subscription.

E3: Thank you for the interview.

oracle.com/sap

Oracle Database Application Continuity for SAP NetWeaver Pilot Program

Oracle Database Application Continuity for SAP NetWeaver

As we wrote in a previous Oracle for SAP Database Update, Oracle Database Application Continuity is a new feature for SAP customers to perform planned downtime on their Oracle database nodes without impacting ongoing activities. It is currently in pilot phase and participation is possible on request.
Active users on SAP NetWeaver application servers connected to an affected database node will not experience any disruption when scheduled database maintenance occurs. System administrators will also not need to reboot or restart application servers due to transaction failures.

More information can be found in SAP Note 3234354, which is currently restricted to pilot program members.

Benefits of Oracle Database Application Continuity

To ensure minimum business disruptions, organizations that require high uptime deploy Oracle Real Application Clusters to ensure that their database systems are always available. As designed, scheduled database maintenance or an outage causes ongoing application sessions to be redirected towards available database nodes.

However, users will encounter errors due to the loss of cached data or in-flight commits if their sessions are active on the database node that is being shut down for maintenance. Moreover, state changes reflected in the database sessions are also lost. This manifests as error messages and SAP short dumps that will vary depending on what users are doing.

With Application Continuity, users will not experience any interruptions during scheduled database or node maintenance due to the automatic replay of interrupted database requests. This happens in the background with no effort from either end users or administrators. The result is an improved user experience and reduced administrative effort when maintenance occurs during operational hours.

Benefits:

- Submitted data, returned data, and variables in the database session are retained.
- Non-transaction database session states are maintained.
- Database and database sessions stay in sync.
- Commit states are correctly recorded.

Overview Concept of Application Continuity

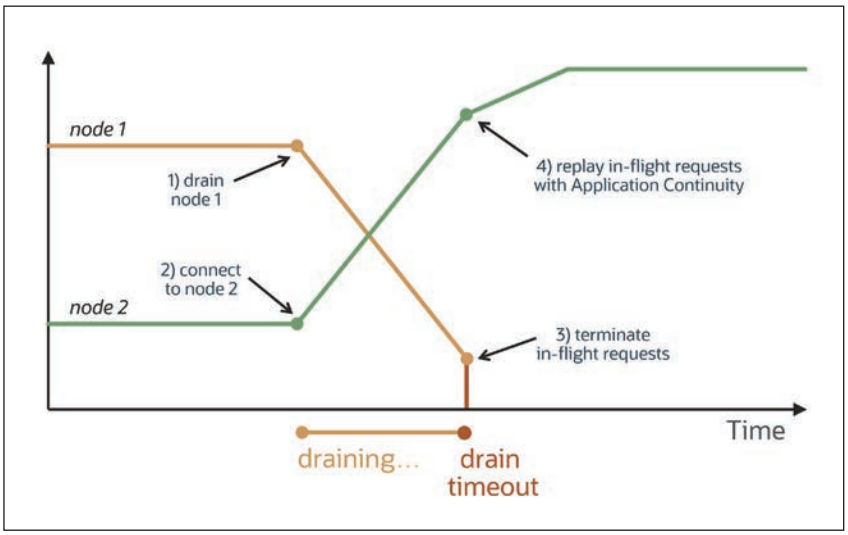

Figure 1: Overview Concept of Oracle Application Continuity

The high-level concept of Application Continuity is based on a "Drain-Connect-Switchover" process.

- Drain work before maintenance. Active sessions allowed to complete within timeout period.
- All new sessions will connect to other available nodes.
- At Drain timeout, remaining sessions are terminated and switched over with automatic replay of in-flight requests on an available node.

A Fast Application Notification (FAN) event is triggered to all registered endpoints when the status of the database service changes and causes notified applications to react quickly. When a planned downtime is initiated, in-flight sessions on the affected node in a Real Application Cluster will begin draining to another node. Draining ends at the end of the configurable drain timeout, and the database node is shut down. At this point, any remaining in-flight sessions will be terminated. With Application Continuity, these sessions are automatically replayed for a seamless transition.

Oracle Application Continuity in Action

As shown in the above diagrams, User 1a and User 1b are connected to DB Node 1 during normal operations (Figure 2). When maintenance is initiated on DB Node 1, new sessions are diverted to DB Node 2 while active sessions are released upon completion. With Application Continuity, ongoing in-flight sessions such as User 1a and User 1b are terminated and replayed automatically on DB Node 2 (Figure 3). User 1a and User 1b continue working with no interruptions or error messages.

With Application Continuity, there is no doubt about the state of a transaction. Crucially, administrators no longer need to restart mid-tier machines and recover from logon storms caused by failed sessions. The

end-user experience is significantly improved by masking planned outages without the need for manual recovery by the application developer.

Planned downtime would be more acceptable to users if the impact on active concurrent users is minimal. With more acceptance, the overall management of the entire infrastructure would be easier with respect to the scheduling and execution of maintenance activities.

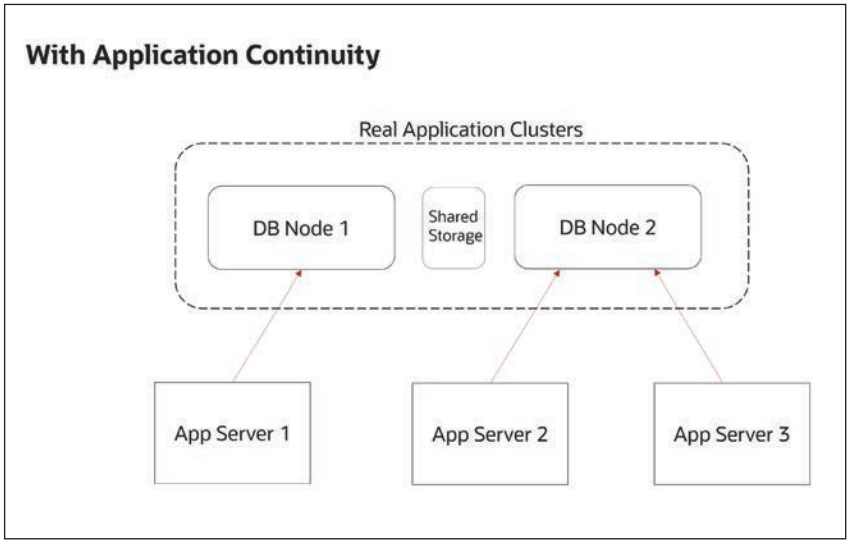

Figure 2: Normal operations

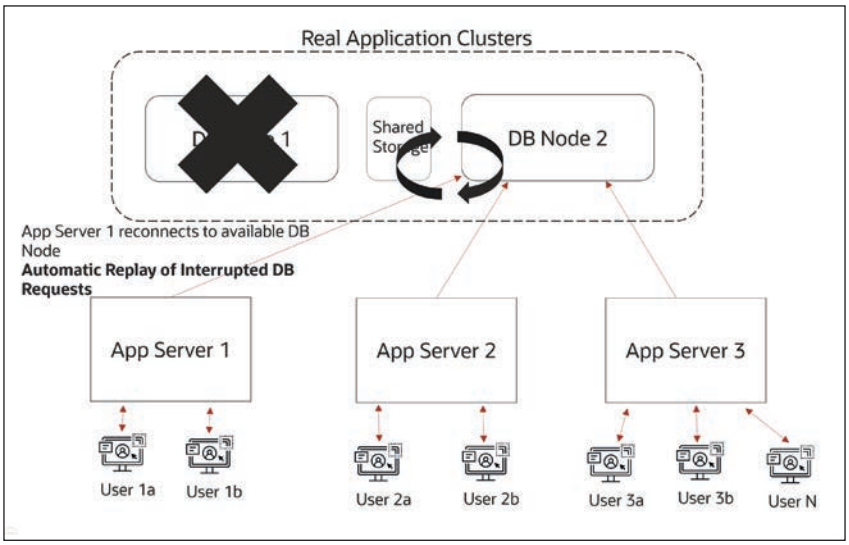

Figure 3: DB Node 2 taken offline for scheduled maintenance

Implementing Application Continuity

The Application Continuity feature is implemented in the DBSL layer of the SAP NetWeaver ABAP Application Server. It is not available in the current regular SAP DBSL layer and will require the installation of a custom DBSL. Additional requirements are the use of database services and some changes to the SAP profiles and configurations for both AS ABAP and AS Java.

To participate in the pilot program, the following hardware and software requirements must be met:

- SAP NetWeaver Application Server 7.50 (both SAP NW AS ABAP and AS Java).
- Special version SAP Kernel 7.53 of SAP DBSL for Oracle.
- Oracle SAP Database Instant Client 19c V4 or newer for SAP NetWeaver.
- Oracle Real Application Clusters 19c (19.15) on Linux (with additional platform support, such as Solaris and AIX, planned after pilot program).

Application Continuity is currently available as part of a pilot phase. Customers may email to frontdesk-walldorf_de@oracle.com to enquire more about this pilot program. Alternatively, they may request participation directly by opening an SAP support message on the **BC-DB-ORA** support queue.

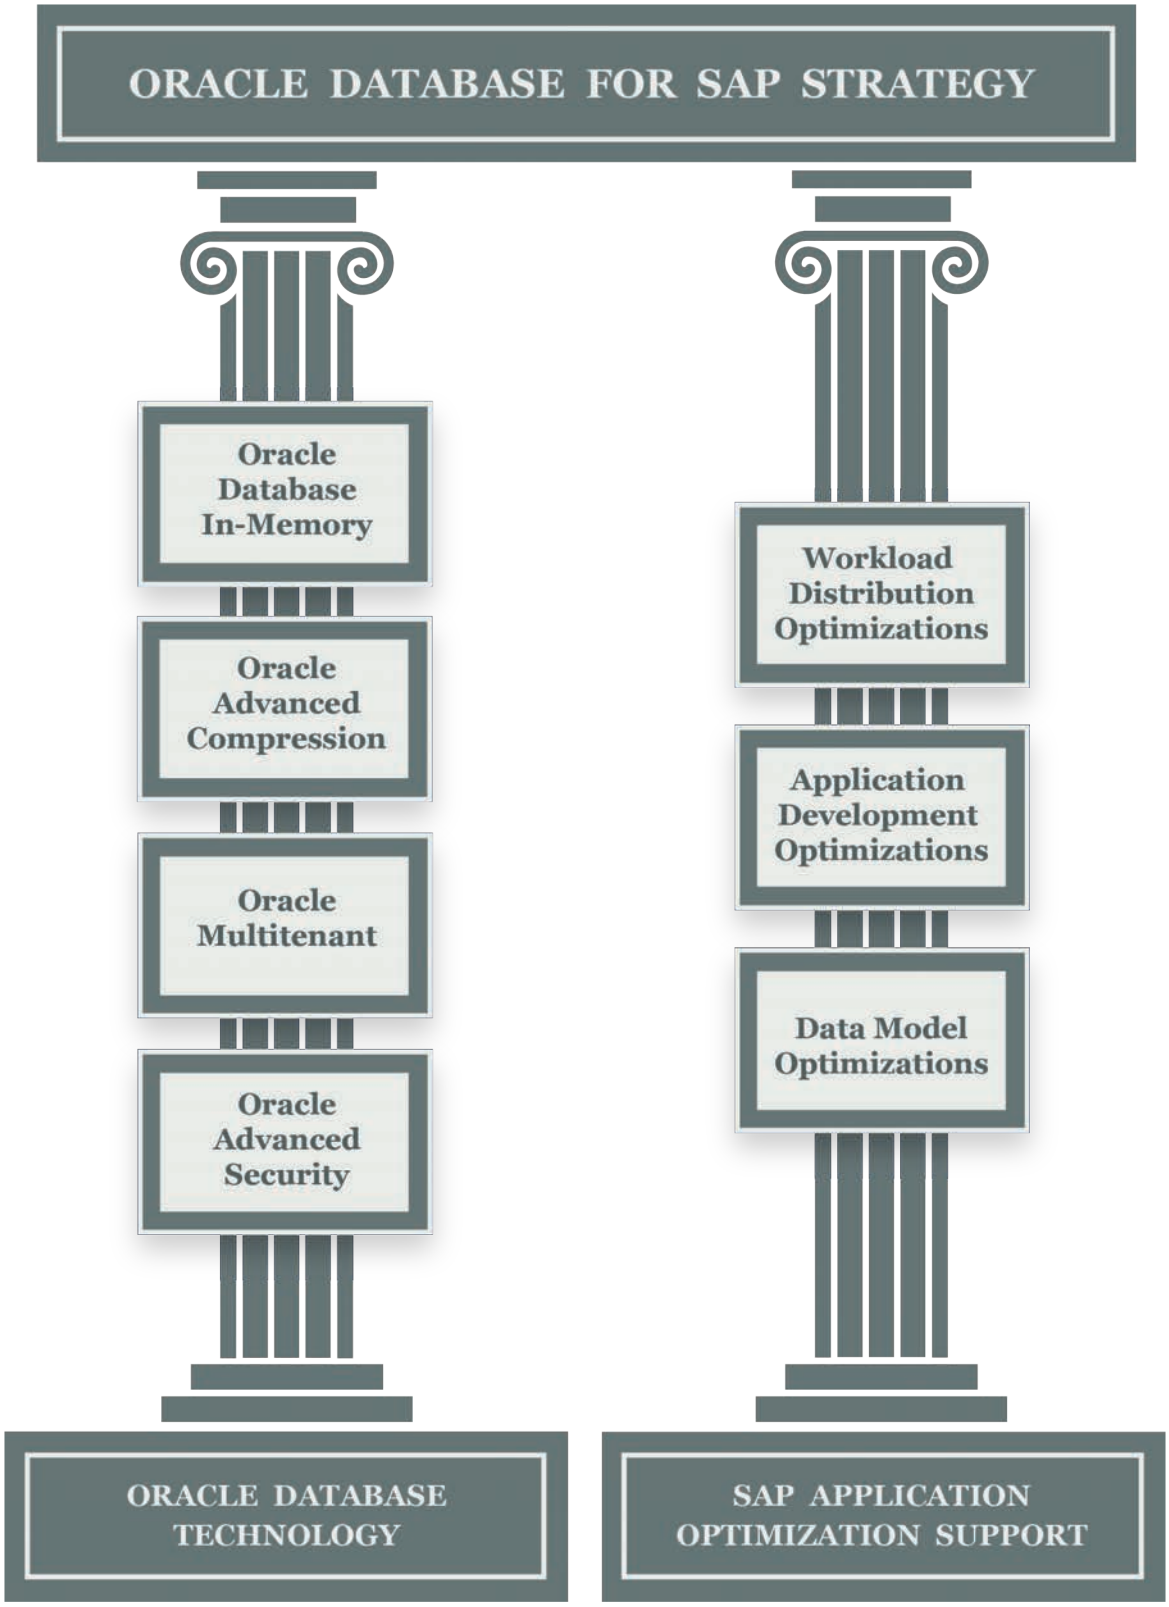
ORACLE DATABASE FOR SAP STRATEGY
Oracle Database In-Memory
Oracle Advanced Compression
Oracle Multitenant
Oracle Advanced Security
Workload Distribution Optimizations
Application Development Optimizations
Data Model Optimizations
ORACLE DATABASE TECHNOLOGY
SAP APPLICATION OPTIMIZATION SUPPORT

Oracle Database for SAP Strategy

Latest Database Technology and Support for SAP Application Optimizations

From the very beginning, the Oracle Database for SAP 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.

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.

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, NetWeaver-based SAP applications (such as R/3 and BW) were 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, if not better. Therefore, the need to integrate SAP application features with the Oracle Database has recently become more visible.

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.

Oracle Database and Database Options

Database Editions

Oracle Database is available in five editions, each suitable for different development and deployment scenarios. However, only Oracle Database Enterprise Edition is certified and supported in SAP environments, as SAP applications are very demanding and cannot be run efficiently without the enterprise computing features provided by Oracle Database Enterprise Edition.

Database Options and Management Packs

In addition, Oracle offers several database options, management packs, and other products that enhance the capabilities of Oracle Database for specific purposes. They extend the power of Oracle Database Enterprise Edition to meet customer- or application-specific requirements in the areas of efficient use of disk space, performance and scalability, high availability, security and compliance, data warehousing, big data, and manageability.

Options and Packs in SAP Environments

This brochure is about database options and management packs for SAP customers. There are some differences between a pure Oracle Database and an Oracle Database for SAP perspective:

- Even if an option is certified, the use of some of its features may not be permitted. As the information provided here is only an overview, not all details can be discussed.

  If in doubt, check SAP Note 105047.

- Due to peculiarities of the SAP data model or application design, an Oracle Database option or management pack may not be optional, but required. E.g., SAP Business Warehouse (BW) on Oracle Database requires Oracle Partitioning.

- From a pure Oracle perspective, an option or pack is licensed separately. However, an Oracle Database Enterprise Edition license bought from SAP (ASFU) already includes some (but not all supported) options and packs without additional payment.

  For details, see SAP Note 740897.

Oracle Database Versions: Support Status and Roadmap

Starting from 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.). There are two

types of Oracle Database releases: Long-term Release and Innovation Release.

Oracle Database Long-term Releases are ideal for use cases that benefit from less frequent upgrades to newer releases. They offer the highest level of stability and the longest length of error correction support. These releases have at least 5 years of Premier Support followed by 3 years of Extended Support. When combined with Extended Support, customers typically have 4 or more years to upgrade from one Long Term Release to the next Long-term Release. – *Oracle Database Long-term Releases are made available (certified) for SAP customers as well.*

In between Oracle Database Long-term Releases, Oracle delivers **Oracle Database Innovation Releases** that include many enhancements and new capabilities which will also be included in the next Long-term Release. Innovation Releases are designed to enable customers to continuously use leading-edge technologies to rapidly develop or deploy new applications or augment existing applications. Support for Innovation Releases includes at least 2 years of Premier Support, but there is no Extended Support. – *Oracle Database Innovation Releases are not certified for SAP.*

Figure 1.1: Oracle Database version support as of July 2024 (Dates are subject to change at any time).

For latest information, see Oracle MOS Note 742060.1 and SAP Note 2606828.

Oracle Database 19c

Oracle Database 19c, certified for SAP since December 2019, is the most current long-term support release, and it is recommended for all SAP on Oracle customers.

Premier Support ends April 30, 2026. Paid Extended Support starts May 01, 2026 and ends April 30, 2027. Error Correction / Patching is available through April 30, 2027 with paid ES. Without paid ES, patching is only available until April 30, 2026.

Oracle Database 23ai

Oracle Database 23ai, the next long-term support release, is planned to be available and to be certified for SAP in CY 2025. This means: The next long-term support release will be supported and certified for SAP before Premier Support for Oracle Database 19c ends.

SAP Application Optimization Support

WORKLOAD DISTRIBUTION OPTIMIZATIONS

Push data-intensive computations to database server
Reduces network traffic, improves application performance.

Push BW FEMS queries to database server
Reduces network traffic, improves application performance.

SAP provides necessary interfaces
E.g. Core Data Services (CDS), FEMS Pushdown

Oracle Database provides base technologies
E.g. support for PL/SQL, stored procedures and functions

APPLICATION DEVELOPMENT OPTIMIZATIONS

Benefit from CDS in SAP standard applications
E.g. ERP 6.0 EHP8, Convergent Invoicing, Banking Services from SAP 9.0

Benefit from CDS in customer-specific applications
Many home-grown applications can benefit significantly

SAP provides development tools
E.g. Eclipse, OData, SAP Fiori

Oracle Database provides relevant technologies
E.g. SQL Macros, which reduce SQL runtime considerably

DATA MODEL OPTIMIZATIONS

Flat Cubes for SAP BW
Optimize SAP BW InfoCubes for in-memory computing.

Declustering/Depooling for SAP ERP
Convert clustered and pooled tables to transparent tables.

SAP provides administration tools
E.g. “Repartitioning Tool” for cube conversion from non-flat to flat.

Oracle Database provides base technology
Oracle Database In-Memory

Oracle Database Support for SAP Application Optimizations

Workload Distribution Optimizations

The Database as Dumb Data Store

What is a database? What can it do? And what can it not do? – The traditional answer to these questions has been that a database is nothing but a dumb data store. It is (we were told) a container that can permanently store data, but that's it. Whenever a user 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.

The Database as Application Tier

30 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 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 in 1992).

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

SAP's Push-down Strategy

When, 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, this behavior 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.

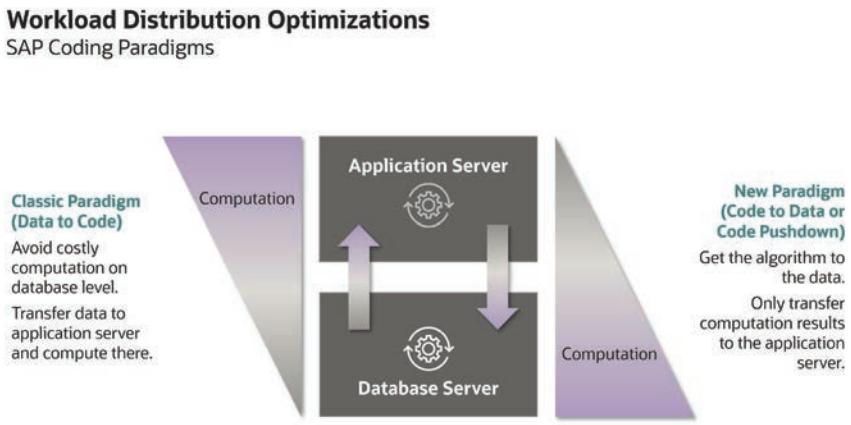

Figure 2.1: SAP coding paradigms

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.

A second example for the same strategy is FEMS Pushdown. FEMS queries can be thought of as a spreadsheet and query conditions that define how to calculate the cell values. FEMS Pushdown, which allows all calculations to be done in the database, can reduce database time, network traffic, and application server time considerably.

For more information on FEMS Pushdown, see SAP Note 2816467.

Application Development Optimizations

A subset of the standard applications developed by SAP uses CDS by default: ERP 6.0 EHP8, Rapid Replenishment Planning, Convergent Invoicing, Banking Services from SAP. However, the benefits of the CDS framework are by no means restricted to SAP standard applications. For customers, home-grown applications are an essential part of their SAP-related portfolio and workload. Many of these apps can significantly benefit from using CDS features.

For details on how to develop your own applications using CDS, see the white paper ABAP Core Data Services. SAP Business Suite – Best Practice Guide.

The white paper can be downloaded from https://tinyurl.com/SAP-ABAP-CDS-on-anyDB

Oracle Database provides the features that are necessary to support this approach. An example is SQL Macros. Originally planned for Oracle Database 21c, this feature has been back ported to Oracle Database 19c. SQL Macros allow developers to factor out common SQL expressions and statements into reusable, parameterized constructs that can be referenced in SQL statements. Unlike PL/SQL functions, SQL Macros are evaluated at parse time, which means that at execution time context

For more information on SQL Macros, see SAP Note 2801989.

switches between SQL and PL/SQL can be avoided and SQL runtime can be reduced considerably.

Data Model Optimizations

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 here) 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. Flat Cubes, which will be explained in this section, utilized 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*.

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.

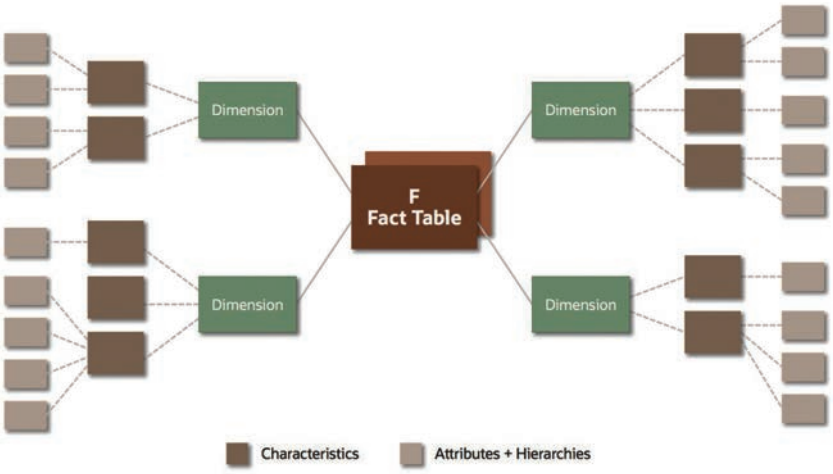

Figure 2.2: Traditional "star" (= extended snowflake) schema

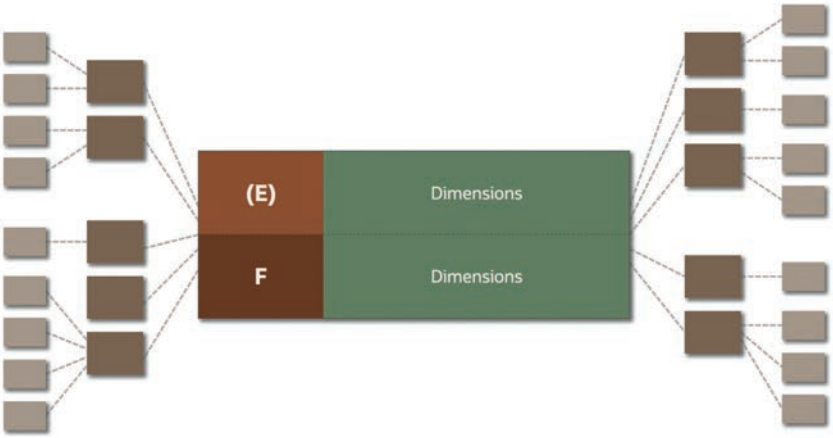

Figure 2.3: New flat cube design

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.

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. The main disadvantages of the flat data model that originally motivated the develop-

ment of the extended snowflake schema have been the disk and memory requirements of storing redundant data. Today, this is no longer a concern, thanks to Oracle's Advanced Compression features, which optimize the storage for data on disk as well as 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.

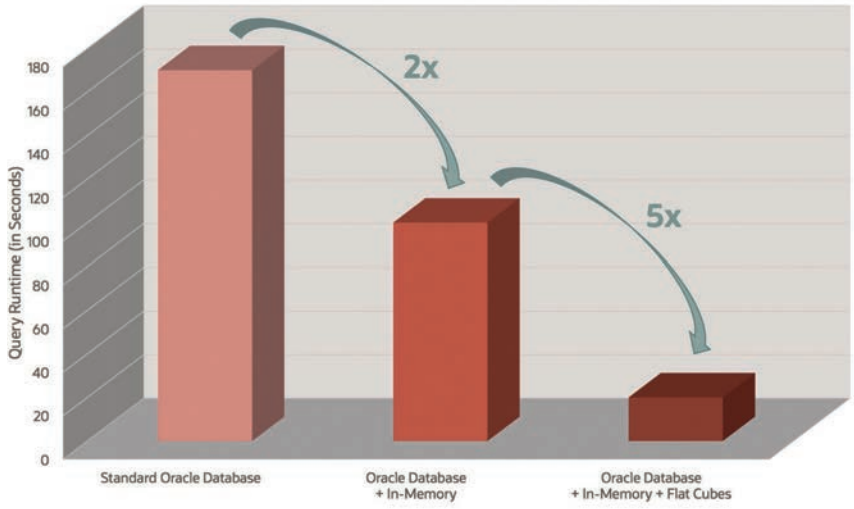

Figure 2.4: Performance gains with Oracle Database In-Memory and Flat Cubes for SAP BW

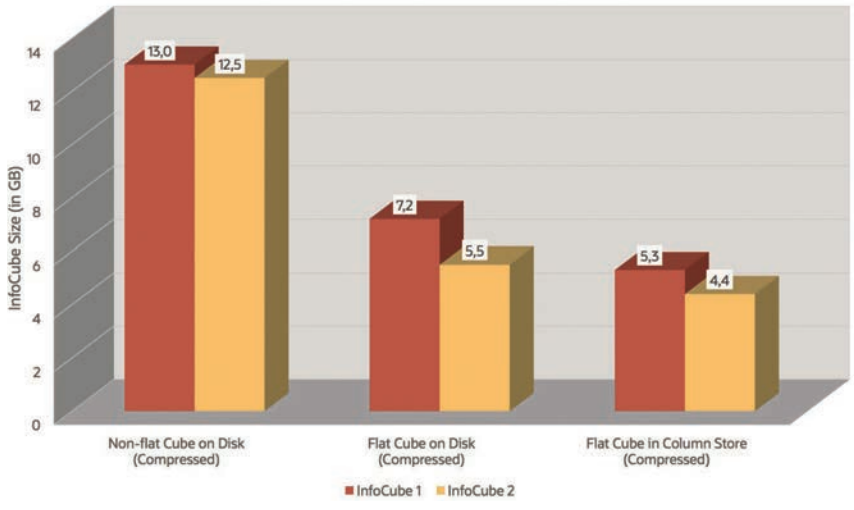

Figure 2.5: Disk space and memory consumption without and with Flat Cubes for SAP BW

Flat InfoCubes: Implementation Tools

As Oracle Database In-Memory provides the same functionality as SAP HANA, SAP and Oracle jointly developed support for Flat InfoCubes on Oracle Database In-Memory. The Flat InfoCubes support includes an extension of the SAP-provided *Partitioning Tool,* which (in addition to its original purpose: table partitioning) may now also be used by system/database administrators to convert traditional InfoCubes to Flat InfoCubes.

No other tool (in particular: no tool provided by Oracle) is needed. If customers want to convert non-flat to flat cubes, the Partitioning Tool is sufficient. *And it is easy to use:* Just select the cube you want to convert. The tool then transforms the tables and determines which of them should be kept in memory.

See figure 3 on page 64

Special cases:

- If you want to convert a large number of InfoCubes, you can use program RSDU_IC_STARFLAT_MASSCONV for mass conversion.

  See SAP Note 2523154.

- Conversion of Semantically Partitioned Objects to a flat format is supported as well. However, only SPOs that consist of InfoCubes can be converted, SPOs consisting of DataStore Objects cannot.

  See SAP Note 2711358.

This may sound like a complete set of technologies and tools, but one element was missing: While customers were able to implement Flat InfoCubes in an existing SAP BW on Oracle system easily, there was no easy way for those customers who wanted to migrate SAP BW on HANA to SAP BW on Oracle. An important step of such a SAP BW migration (a.k.a. heterogeneous system copy) is the report SMIGR_CREATE_DDL. It is run in the source system, and it creates DDL (in particular: CREATE TABLE) statements for non-standard objects. The output is then used to build the target system.

If a BW system is migrated from AnyDB/Oracle to HANA, traditional InfoCubes are by default converted to Flat InfoCubes. SMIGR_CREATE_DDL knew very well that these new InfoCubes were non-standard objects, but it was neither aware of Oracle's support for Flat InfoCubes, nor did it know the SQL syntax used to build Flat InfoCubes in the Oracle Database.

The missing information was recently added, and as of December 2020, the enhanced SMIGR_CREATE_DDL report is generally available.

Details can be found in SAP Note 2948714.

Oracle Database In-Memory

INNOVATION FOR EFFICIENCY AND FLEXIBILITY

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

Ask New Questions
Queries that took too long in the past

Implement Simplified SAP Data Models
E.g. Flat Cubes

Simplify Customizations
E.g. get rid of user-defined indexes

PLATFORM CONTINUITY

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

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

No Data Migrations or Reorganizations
Storage location and format unchanged

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

EASY TO UNDERSTAND AND TO DEPLOY

Define Size of In-Memory Column Store
Oracle Database Server initialization parameter

Select Tables to be Populated into In-Memory Column Store
Any subset of existing tables, no minimum or maximum

SAP-provided Configuration Tool for Business Warehouse
Cube conversion from non-flat to flat

SAP-provided Migration Tool for Business Warehouse
Flat cube migration from non-Oracle to Oracle

Oracle Database In-Memory

Oracle Database 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 it makes sense. In all other cases, data are stored and handled as it always has been.

Memory: The New Dual-Format Architecture

Oracle Database has traditionally stored data in a *row format.* This format is ideal for online transaction (OLTP) systems, as it allows quick access to all columns in a record. A *column format* database stores each of the attributes about a transaction or record in a separate column structure. This format is ideal for analytics, as it allows for faster data retrieval when only a few columns are selected but the query accesses a large portion of the data set.

But what happens, when your system is characterized by a mixed workload? Up until now you have been forced to pick just one format and suffer the trade-off of either sub-optimal OLTP or sub-optimal analytics performance. The only way to optimize for both OLTP and analytics had been to copy data from OLTP systems to analytic systems using complex ETL processes that add a great deal of expense and latency.

Oracle Database In-Memory optimizes both 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. This architecture eliminates the trade-off 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 transactions 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 today.

Disk: Nothing Has Changed

The new column format is a pure in-memory format. Tables are stored on disk using Oracle's existing row-based (or – on Engineered Systems – hybrid columnar) formats. Since there is no persistent columnar storage

Challenge: In more and more systems, meeting analytics performance requirements turns out to be a challenge. This is true for long-running queries in SAP BW. However, it can also happen in OLTP systems, e.g. if a very flexible implementation of operational planning/reporting allows users to create many, slightly different query variants.

Value Proposition: Oracle Database In-Memory allows administrators to dedicate a certain amount of database server memory to the Column Store – a memory structure that stores data in column format instead of in row format. Setup of the Column Store is fast and easy. Having data available in column format can improve query performance substantially.

Certification/Support: Oracle Database In-Memory is certified for all SAP NetWeaver applications.

Versions: Oracle Database 12c and higher

Implementation: For an overview and pointers to more detailed documents, see SAP Note 2178980.

format, there are no additional storage costs or storage synchronization issues. Nor is there a need to modify the database. Oracle Database In-Memory can be implemented without a database migration or a table reorganization.

As a result, the new Oracle Database 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). These features work exactly the same way with and without Oracle Database In-Memory

.

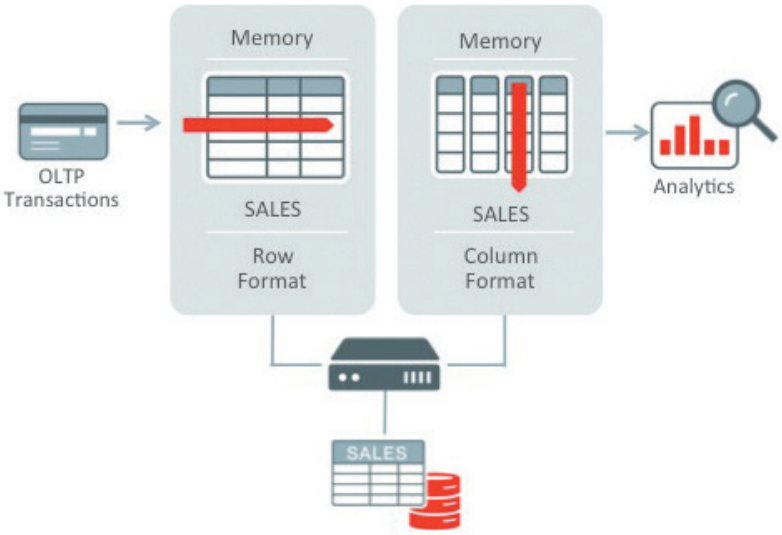

Figure 3.1: Oracle Database In-Memory – Dual memory, single disk format

Easy to Implement and Manage

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 class 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

As an example, see the Bosch customer story on page 52.

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 round: 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 <table_name> 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.

Fine-Grained Control

An easy start based on intelligent defaults for typical situations – this is what Oracle customers expect. In addition Oracle customers expect mechanisms, which allow for fine-grained control and tuning. Oracle Database In-Memory provides such mechanisms. Examples are:

- 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. Therefore, administrators may want to restrict the population process to the data really needed by DSS queries. Table partitioning allows them to make this happen. If the table is partitioned in a useful way (e.g. by month), this internal structure can be used to define a horizontal subset of the table data to be kept in the in-memory column store.
- The Oracle Database has been optimized and tuned for decades to scale-up on SMP servers. In addition to being able to scale up, Oracle Database In-Memory can also scale out to very high memory and CPU capacities by using all of the memory and processors in a cluster of servers (RAC). In such environments, all objects populated into memory will by default be distributed across all in-memory column stores in the cluster. On Oracle Engineered Systems, objects can also be duplicated. This means that an object (or part of an object, e.g. a partition) populated into the in-memory column store will have a mirrored copy placed on one of the other nodes in the RAC cluster. Duplicating data provides in-memory fault tolerance as it ensures data is still accessible via the in-memory column store even if a node goes down or is taken down for maintenance.

Deployment Flexibility

HARDWARE AND OPERATING SYSTEM PLATFORM SUPPORT

Generic Servers
Unix, Linux, Windows

Oracle Engineered Systems
Oracle Exadata Database Machine

Oracle Cloud
Oracle Cloud Infrastructure, Exadata Cloud Service, Exadata Cloud@Customer

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)

DATABASE MIGRATION SERVICES

Non-Oracle to Oracle
Non-Oracle database to Oracle Database migration

Oracle to Oracle
Different versions; different sets of features; storage optimization

Oracle to Oracle close to Zero Downtime
Minimize downtime to almost zero

On-Premise to Cloud
Part of a migration to cloud

Deployment Flexibility

Implementation and Migration 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 is 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 three 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.
- *Oracle Exadata Cloud@Customer* is the simplest way to move an organization‘s business-critical Oracle Database workloads to the cloud. It simultaneously runs Oracle Exadata Database Service and the fully managed Oracle Database Service inside customers‘ data centers and behind their firewalls to help meet strict data residency and security requirements.

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.

Oracle Multitenant is an Oracle Database option that helps customers reduce IT costs by simplifying consolidation, provisioning, upgrades, and more. It relies on an 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.

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.

Data Guard: Data Guard can be combined with all previously mentioned architectures. While RAC provides high availability by multiplying the number of Oracle instances, 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.

Oracle Transition Service for SAP Migration

Through automation, advanced support tools, and more than 15 years of Oracle Database transition experience, Oracle Transition Service delivers transition planning, validation, and execution services that allow you to transition your SAP database faster, with lower costs and less downtime. Oracle experts assess your SAP environment, guide you in choosing from multiple transition approaches, resolve potential issues, perform test runs to reduce risk, and move your database through a secure online gateway. Oracle's efficient process, based on the latest industry best practices, allows you to transition your SAP database in a matter of days, not the weeks or months typically required.

Oracle Transition Service is delivered by Oracle Advanced Services (ACS), which, as of July 2023, is part of Oracle Customer Success Services (CSS).

Architecture Option: Oracle Multitenant

Consolidation Approaches

Large enterprises may use hundreds or thousands of databases. Often these databases run on different platforms on multiple physical servers. A database may use only a fraction of the server hardware capacity. This is an expensive approach which fails to maximize the usage of both the hardware and human resources.

A typical response to the management problem is to place multiple databases on each server (either as direct installs or using virtual machines). The problem is that the multiple database instances do not share background processes, system and process memory, or Oracle metadata. Another response is to logically separate the data into schemas (schema consolidation). The problem is that these virtual entities are difficult to manage, secure, and transport.

Oracle Multitenant Architecture

Oracle Database Multitenant is based on an approach called database consolidation. It relies on an architecture that allows one single Container Database (CDB) to hold many Pluggable Databases (PDBs).

An existing database can simply be "plugged into" a CDB. At any time, then, it can be unplugged and plugged into another CDB. Unplug/plug is even supported across Oracle Database software versions.

Challenge: Many SAP landscapes consist of a few large and a considerable number of small or very small systems. However, the existence of many small SAP systems based 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.

Value Proposition: Oracle Database Multitenant introduces an architecture that enables customers to easily consolidate multiple databases, without changing their applications. This architecture delivers all the bene-

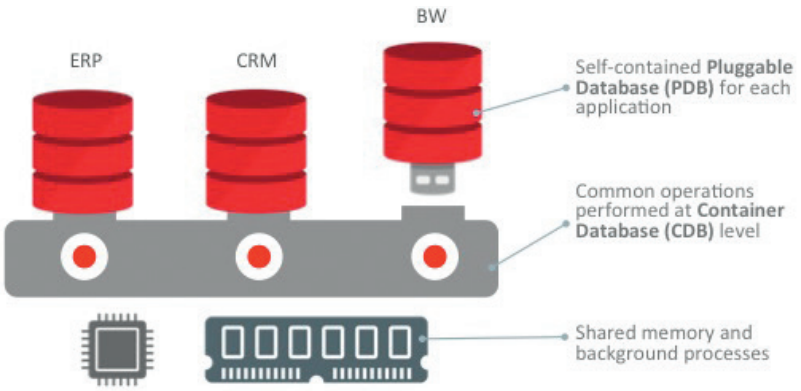

Figure 4.1: Oracle Multitenant – Architecture for consolidating databases and simplifying operations

From the point of view of the client application connecting to the database server via Oracle Net, the PDB is the database. A PDB is fully compatible with a non-CDB – a rule also known as the PDB/non-CDB compatibility guarantee.

Resource Utilization and Resource Management

The many PDBs in a single CDB share its memory and background processes. This enables consolidation of many more databases compared to the traditional single-instance architecture, offering similar benefits to schema-based consolidation but with none of the major application changes required by that approach.

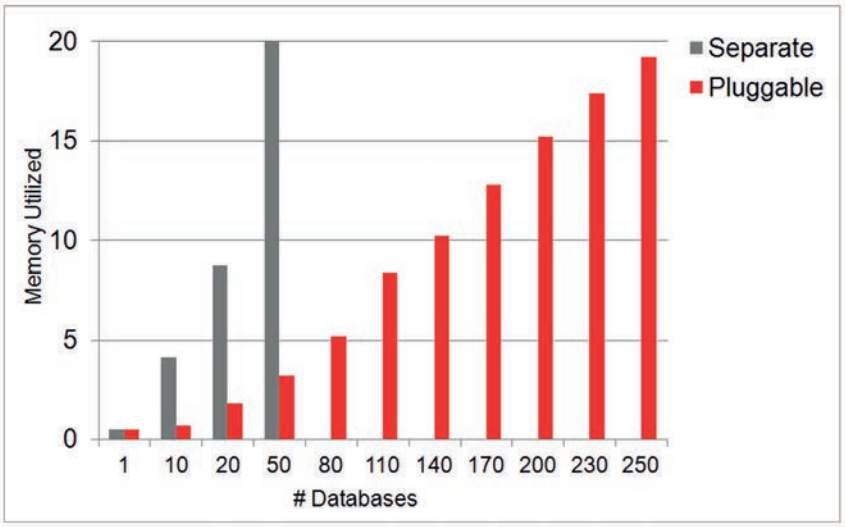

Figure 4.2: Oracle Multitenant – Resource utilization benchmark results

- *Only 3GB of memory vs. 20GB memory used for 50 databases.*
- *Multitenant architecture scaled to over 250 databases while separate database instances maxed at 50.*

Horizontal partitioning of the Oracle data dictionary (a conceptual partitioning of managing many databases as one, yet retains the isolation and resource prioritization of separate databases. Oracle Multitenant reduces resource consumption by separating "container" and "pluggable" databases. It simplifies administration by moving standard operations to the "container database" level.

Certification/Support: Oracle Multitenant can be used in SAP environments.

Version: Available as of Oracle Database 12c. Mandatory as of Oracle Database 23ai.

Implementation: For more information, please see SAP Notes 2336881, 2335850, and 2333995.

tioning, not a physical table partitioning) removes the need to store and manage system-wide metadata in every single database. The “lower” half (implemented in the CDB) holds the system-wide metadata – and nothing else, while the “upper” halves (implemented in the PDBs) hold application-specific metadata – and nothing else.

Creating pluggable databases, moving pluggable databases between containers, and cloning pluggable databases are done with new SQL commands and take just seconds. When the underlying file system supports thin provisioning, many terabytes can be cloned almost instantaneously.

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

Manage Many Databases as One

By consolidating existing databases as pluggable databases, administrators can manage many databases as one. Benefits include:

- The investment of time and effort to patch one CDB results in patching all of its many PDBs. To upgrade all PDBs hosted in a CDB, simply upgrade the CDB, and 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. In other words, all PDBs consolidated into a container will be backed up as one, and 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 standby configuration at the CDB level, to replicate all PDBs consolidated in that container.

Oracle Multitenant for SAP

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.

Architecture Option: Real Application Clusters (RAC)

In order to guarantee high availability of the database server, traditionally a failover cluster had been implemented. Such a solution, however, had at least two disadvantages:

- A failover cluster relies on the concept that, at any given moment, only one database instance running on one machine can be active. The other machine (most probably an expensive server, too) is always idle.
- When a problem on the primary machine is detected, an Oracle Database Server instance needs to be started on the secondary machine. In this particular situation, startup can take up to 30 minutes – which means: up to 30 minutes unplanned downtime.

In the Real Application Clusters (RAC) architecture, multiple Oracle instances are up and running at the same time. Therefore no restart is required. If one of the RAC servers fails, the other instances can take over. A reconnect of the affected users is a matter of seconds, not of minutes.

Benefits

To summarize the benefits in a few words: Oracle Real Application Clusters combines workload distribution, scalability, high availability, better manageability, and cost savings.

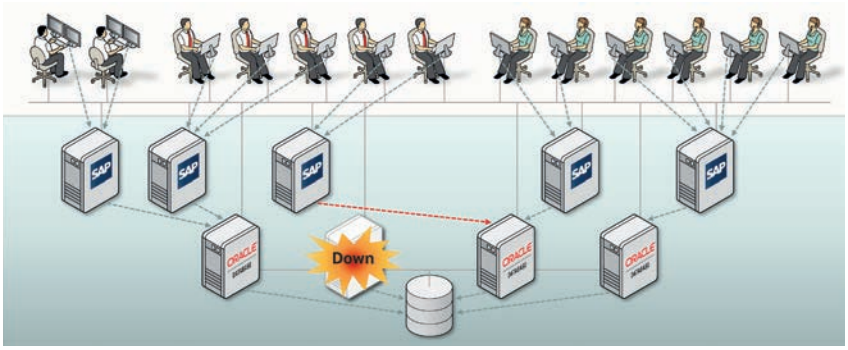

Figure 4.3: Real Application Clusters (RAC) for scale-out and immediate (instance) fail-over

Scalability: Using RAC, the scale-out approach, which is always supported on the SAP Application Server level, can be implemented on the database level too. In this example, 5 SAP Application Server instances, running on 5 different machines, are connected to 4 Oracle Database Server instances, running on 4 different machines.
High Availability: If one of the Oracle instances goes down, the affected SAP instance(s) is/are automatically reconnected to one of the available Oracle instances. After this operation users can continue their work. The fail-over occurs within seconds.

Cost Savings: Multiple small servers are much cheaper than one single big server. In addition, there is no need to buy them all at the same time. They can be bought and added as needed.

Challenge: When the workload on a database server increases (due to new application versions, additional applications, or more users), the traditional solution is to replace the existing server with a bigger one (scale-up). Big servers, however, are very expensive.

Value Proposition: Real Application Clusters (RAC) allows multiple instances to be up and running and to access the same database at the same time.
As these instances can and in most cases do run on different machines, customers have the option to implement a scale-out approach: 2 or more small servers can handle the same workload as one big server. However, they are much cheaper. And they can be added (or removed) as needed.

Certification/Support: Oracle Real Application Clusters is certified for all SAP NetWeaver applications.

Versions: Oracle Database 11g and higher

Implementation: Customers can use any general-purpose machine (Unix, Linux, or Windows) certified by SAP to build a RAC system. In addition, Oracle offers Engineered Systems, which make the implementation considerably easier. Oracle Grid Infrastructure provides a set of base technologies that simplify implementation and help save money.

Oracle Grid Infrastructure and SAPCTL

Oracle Grid Infrastructure provides the base technologies that are required to enable RAC. It can be divided into two main components:

- In order to let multiple Oracle instances access the database files at the same time, a cluster file system is needed. Therefore Oracle provides *Oracle Automatic Storage Management (ASM).* Unlike other, third-party cluster file systems, it is optimized for Oracle Database files, and it is available for free.
- *Oracle Clusterware* is the cross-platform cluster software required to run the RAC option for Oracle Database. It enables the nodes to communicate with each other, allowing them to form a cluster of nodes which behaves as a single logical server. Similar to Oracle ASM, which eliminates the need for a third-party cluster file system, Oracle Clusterware eliminates the need for third-party cluster management software.

Oracle Clusterware can provide high availability and resource management for SAP resources just as it does for Oracle resources. Therefore Oracle/SAP Development has created an Oracle Clusterware tool, *SAP Control (SAPCTL),* to enable customers to easily manage SAP high availability resources.

Architecture Option: Data Guard / Active Data Guard

Data Guard can provide both zero data loss protection and near-immediate restoration of service should a production database become unrecoverable for any reason. This is accomplished using the combination of Data Guard synchronous redo transport and a replication-aware apply process at the standby database

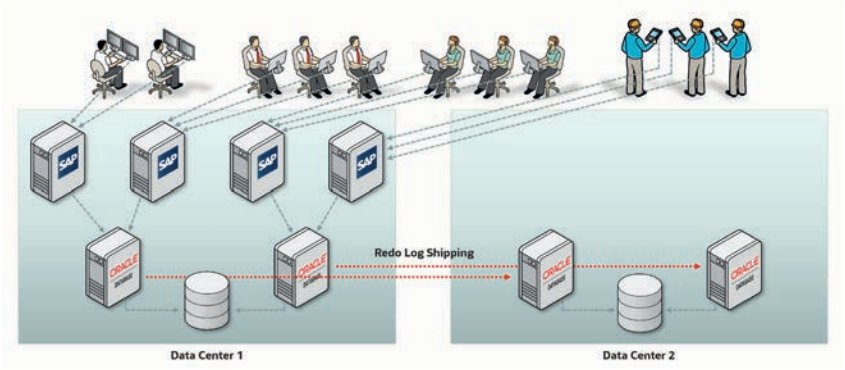

Figure 4.4: Oracle Data Guard – Redo log shipping from production database (Data Center 1) to standby database (Data Center 2) during normal operation

Data Guard can provide both zero data loss protection and near-immediate restoration of service should a production database become unrecoverable for any reason. This is accomplished using the combination of Data Guard synchronous redo transport and a replication-aware apply process at the standby database.

Challenge: RAC provides high availability by multiplying the number of Oracle instances. Such high availability, however, is restricted to the instance level. Even in a RAC-based system, the database remains a single point of failure. This means that DBA errors, data corruption, server or data center failures can make the whole system unavailable.

Value Proposition: 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.

Please note that Data Guard is included in Oracle Database Enterprise Edition. It is not an option. However, Active Data Guard is an option. It offers additional features such as Automatic Block Repair and Fast Incremental Backup.

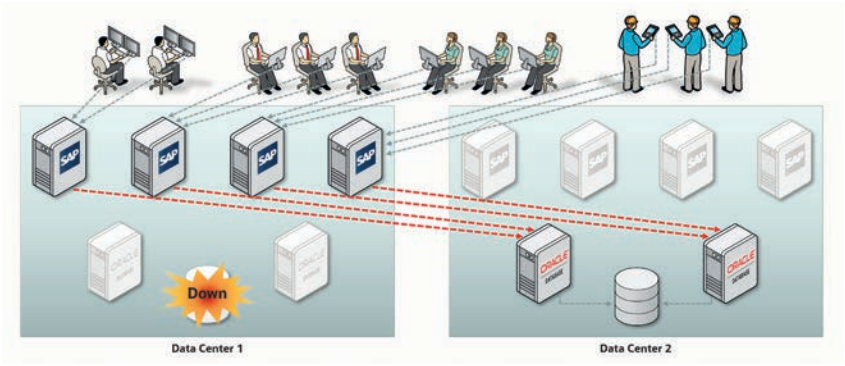

Figure 4.5: Oracle Data Guard – Switch to standby database after failure of production database

- *Option 1: Users remain connected to SAP Application Servers in Data Center 1, SAP Application Server instances reconnect to standby database in Data Center 2*
- *Option 2: After startup of SAP Application Servers in Data Center 2, users reconnect to SAP Application Server instances in Data Center 2, which are connected to standby database in Data Center 2*

The impact that any synchronous replication method can have on database performance, however, often makes it impractical to implement zero data loss protection when large distances separate the primary and replica database(s). Rather than impact database performance, many enterprises will compromise on data protection by implementing asynchronous replication and accept that an unrecoverable outage will result in varying degrees of data loss.

Active Data Guard Far Sync eliminates the need to compromise by extending zero data loss protection to a replica database located at any distance from the primary database. Far Sync provides zero data loss protection for a production database by maintaining a synchronized standby database located at any distance from the primary location, and can do so without performance impact and with minimal cost or complexity. A new type of Data Guard destination called a far sync instance receives changes synchronously from a primary database and forwards them asynchronously to a remote standby. Production can be quickly failed over, manually or automatically, to the remote standby database with zero data loss.

A far sync instance is a light-weight entity that manages only a control file and log files. It requires a fraction of the CPU, memory, and I/O of a standby database. It does not have user data files, nor does it run recovery. Its only purpose is to transparently offload a primary database of serving remote destinations. A far sync instance can save network bandwidth by performing transport compression using Oracle Advanced Compression.

Take for example an existing asynchronous Data Guard configuration with a primary in Boston and a standby in San Francisco. Upgrade to zero data loss simply by using Active Data Guard to deploy a far sync instance within synchronous replication distance of Boston (less than 150 miles). There is no disruption to the existing environment nor is there any requirement for proprietary storage, specialized networking, more database licenses, or complex management.

Certification/Support: Oracle Data Guard is certified for all SAP NetWeaver applications. However, only physical standby databases are supported, logical standby databases are not. Oracle Active Data Guard is certified for all SAP NetWeaver applications. However, Real-Time Query is not possible in SAP environments, because even report generation is not a read-only operation.

Versions: Oracle Database 11g and higher

Implementation: Standard Oracle setup procedures apply. In the white paper "Oracle Standby Database" SAP describes BR*Tools support.

Very Large Database Support

TABLE AND INDEX COMPRESSION

Advanced Compression
Reduce disk space needed for database by 50% or more

Advanced Compression
Improve database server performance by reducing I/O and storing more data in buffer cache

Hybrid Columnar Compression
Use even stronger compression for cold (historical) data

In-Memory Compression
Reduce memory needed for column store and improve analytics performance

TABLE AND INDEX PARTITIONING

Improve Database Server Performance
Reduce runtime of batch jobs (e.g. monthly reports)

Reduce Database Size
Reclaim fragmented unused disk space

Keep Database Size Stable
Avoid fragmentation, e.g. as a consequence of SAP Archiving

Improve Database Manageability
Foundation for Information Lifecycle Management (ILM)

WORKLOAD DISTRIBUTION

Support for SAP's Push-down Strategy
Push down data-intensive application logic to database server

Real Application Clusters
Distribute same workload across multiple instances/servers

Real Application Clusters
Run different workloads on different instances/servers

Data Guard
Run part of the workload on standby database

Very Large Database Support

Compression

Advanced Compression

Figure 5.1 shows a typical Oracle Database which forms an integral part of an SAP (in this case: SAP ERP) system. Approximately one third of the allocated disk space is used for indexes (red), and two thirds contain table data (blue). Table data in turn can be divided into structured data (organized in columns) and unstructured data (PDF or image files, table data "clustered" by SAP).

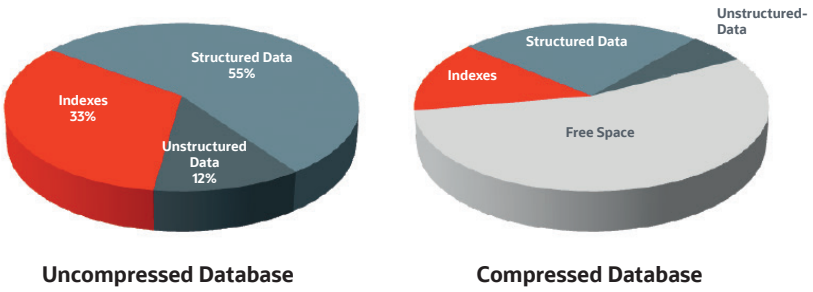

Figure 5.1: Oracle Database Index Key Compression and Advanced Compression (OLTP Compression, SecureFiles Compression)

Oracle Database can compress all three types of data:

- **Index Key Compression** is available for indexes. Index-Organized Tables (IOTs) can be compressed as well. Both features are included in Oracle Database Enterprise Edition, and do not require Advanced Compression.
- **OLTP Compression**, a major feature of Advanced Compression, can be used to compress structured table data. It is not restricted to OLTP systems, but can be implemented in SAP BW systems as well.
- **SecureFiles Compression** (an Advanced Compression feature, too) can be used to compress unstructured table data.

If all features are implemented and all appropriate database objects are compressed, customers have seen 55% disk space savings on average. (This assumes a completely reorganized database. If a fragmented database has not been reorganized, the effects of reorganization and compression are combined, in which case customers have seen up to 80% disk space savings).

The maximum disk space savings which can be achieved depend on the characteristics of the data, and the data characteristics depend on the SAP applications used. Usually SAP BW data can be compressed more efficiently than SAP ERP data, and SAP CRM data permit even higher disk

Challenge: In more and more cases today the size and the expected future growth of the database becomes a problem. Aspects of this problem include: Storage cost, performance guarantees (SLAs), cloning and backup of database files within a reasonable time.

Value Proposition: Oracle Advanced Compression uses a different format for storing table data. Together with other compression technologies, which come with Oracle Database Enterprise Edition (e.g. Index Key Compression), it helps reduce the database size by 50% or more. This is the essential benefit in the sense that this is the effect Advanced Compression is designed for. The benefit from a smaller source database footprint is that the creation of backups and other copies will require less time.

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

Certification/Support: Oracle Advanced Compression is certified for all SAP NetWeaver applications. Implementations are supported by SAP.
Versions: Oracle Database 11g and higherImplementation: Oracle Advanced Compression can be implemented easily in SAP environments, as SAP provides the tool BRSPACE, which is aware

space savings.

Oracle Database Advanced Compression offers more than OLTP and SecureFiles compression. Backup files created by RMAN and export files created by Data Pump can be compressed substantially, even if tables and indexes in the production database are already compressed. Additionally, redo log data can be compressed as well before being shipped from the production to a standby database.

of all SAP-specific requirements. For details, check SAP Note 1431296.

See the Data Guard section on page 27.

Heat Map and Automatic Data Optimization

Based on the information provided so far, it might seem that compression simply reduces the disk space needed, and has nothing to do with database structure. This is an illusion. We need to distinguish between tables that benefit and tables that do not benefit from compression (if this were not the case, compression could be made the default), i.e. between tables that should and tables that should not be compressed.

Nevertheless, this is still a very basic and inflexible distinction. Take, e.g., an SAP BW table that is used for data load. On the one hand, such a table could be compressed, because for most of the time it is accessed in read-only mode. On the other hand, it should not be compressed, because this would slow down the load operations considerably.

Heat Map and *Automatic Data Optimization* allow you to introduce a new differentiating parameter: If a table or partition should be compressed, when would you like it to be compressed? In other words: You can specify that data should be loaded today and (automatically) compressed tomorrow. Heat Map automatically tracks modification and query time-stamps at the row and segment levels, providing detailed insights 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.

Deferred Compression and Information Lifecycle Management

Oracle Database 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 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.

Hybrid Columnar Compression (HCC)

In addition to algorithms for compression of structured and unstructured data, Oracle Exadata supports 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.

Customers running Oracle Database on Oracle Engineered Systems such as Exadata 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 (for 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;
```

Table and Index Partitioning

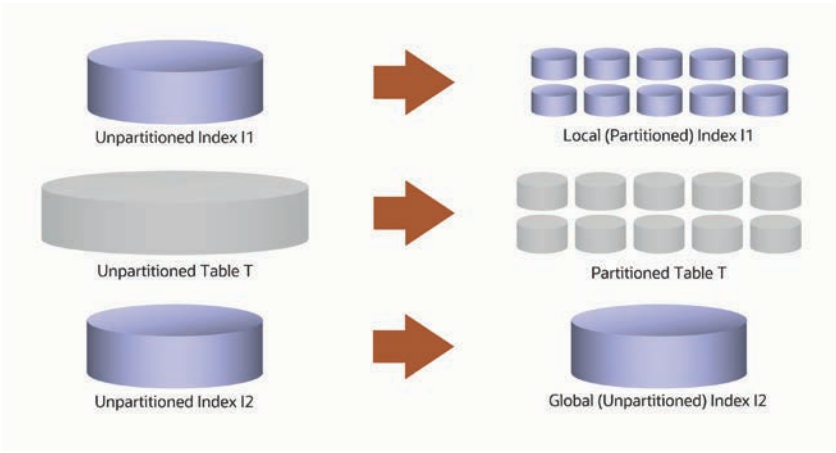

Figure 5.2: Oracle table and index partitioning options

Partitioning enables tables and indexes to be subdivided into individual smaller pieces. Each piece of the database object is called a partition. A partition has its own name, and may optionally have its own storage characteristics. From the perspective of a database administrator, a partitioned object has multiple pieces that can be managed either collectively or individually. This gives the administrator considerable flexibility in managing a partitioned object. However, from the perspective of the application, a partitioned table is identical to a non-partitioned table; no modifications are necessary when accessing a partitioned table using SQL DML commands. Logically, it is still only one table and any application can access this one table as they do for a non-partitioned table.

Figure 5.3 explains why storing related data as closely together as possible makes such a difference. It shows database blocks filled with records. The different colors represent criteria such as different months or different locations. And we assume that the applications accessing those data in most cases want to retrieve all records having the same color.

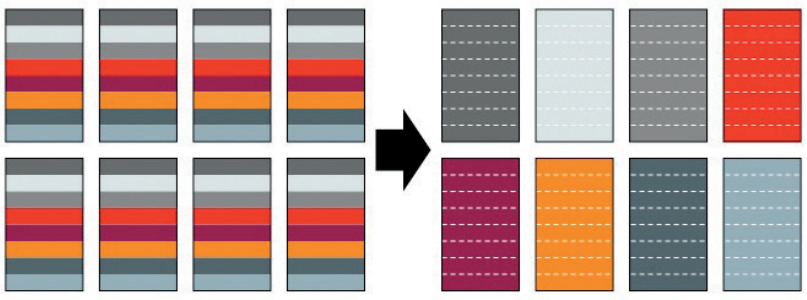

Figure 5.3: Table Partitioning – Physically store subsets of related data as closely together as possible

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

(a) Single queries or complex batch jobs accessing a certain subset of the table data need too much time to complete.

(b) 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.

(c) Data archiving results in heavily fragmented databases.

Value Proposition: 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:

(a) 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.

(b) 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.

(c) The data archiving strategy can be based on the partition structure, and this can avoid disk space fragmentation.

Certification/Support: Oracle Partitioning is certified for all SAP NetWeaver applications.

Under these circumstances, the situation shown on the left hand side is the worst that can be imagined: Each database block contains one record of each color. Or, to put it differently, all subsets of records having the same color are distributed over all blocks. This is the worst possible situation from an I/O perspective (a query looking for all records of a certain color needs to read 8 blocks), from a memory perspective (even if all users work with records of the same color, all 8 blocks need to be cached completely in database memory), therefore from a performance perspective (too much I/O means unsatisfactory performance), from a database administration perspective (subsets of records having the same color cannot be managed independently), and from an ILM perspective ("hot" and "cold" data cannot be separated).

The situation shown on the right hand side, in contrast, is optimal (again: under the circumstances described above): All records having the same color are now stored together in one and the same database block. This is optimal from an I/O perspective (a query looking for all records of a certain color needs to read only 1 block), from a memory perspective (if all users work with records of the same color, only 1 block needs to be cached in database memory), therefore from a performance perspective (significantly less I/O means significantly better performance), from a database administration perspective (subsets of records having the same color can be managed independently), and from an ILM perspective ("hot" and "cold" data can be separated).

Now, multiply the number of records and blocks shown on figure 5.2. Then a partition is the subset of all blocks containing records of the same color.

Workload Distribution

Other options focus on workload distribution. A prominent example is Real Application Clusters (RAC), which allows customers to split the system workload and let many servers, running at least as many Oracle instances, handle the workload. It is up to the customer to decide whether all instances should handle the same type (or mixture) of workload(s) or different instances should be responsible for different types of workload (e.g. interactive transactions vs. batch jobs).

Versions: Oracle Database 11g and higher

Implementation: Partitioning (range partitioning) is configured and used by default in SAP BW on Oracle. In SAP OLTP systems, it can either be implemented by customers themselves or by Oracle Advanced Customer Service (ACS), which, as of July 2023, is part of Oracle Customer Success Services (CSS).

Database Security

ADVANCED SECURITY: ENCRYPTING DATA IN TRANSIT AND AT REST

Network Encryption
Encrypt data travelling between SAP Application Server and Oracle Database Server

Transparent Data Encryption
Encrypt data in production database files

Backup Set Encryption
Encrypt data in backup copies

SAP Notes
- 973450
- 1324684
- 2591575

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.

SAP Notes
- 1355140
- 1502374

AUDITING: MONITORING USER DATABASE ACTIONS

Traditional Auditing
Standard auditing mechanism provided by Oracle Database

Product-specific Auditing
Additional auditing functionality provided by database options (e. g. Database Vault)

Unified Auditing
Capture audit records from a variety of sources in one single audit trail

SAP Note
- 2379567

Database Security

Encrypting Data in Transit and at Rest

Protecting Data in Transit: Oracle Network Encryption

In SAP environments, users do not directly connect to the Oracle Database server. They connect to an SAP application server instance, and the SAP application server instance in turn connects to the Oracle Database server. So, in this case, the application server instance is the Oracle client, and Oracle's network encryption encrypts all data traveling between application server and database server.

Oracle's network encryption requires Oracle software, which is not installed on end-user devices. Therefore other technologies or products must be used to protect the communication between SAP users and SAP application servers.

However, people do not only attempt to read data in transit, they can also try to intercept and modify them. Therefore, in addition to network encryption, Oracle also supports crypto-checksumming to ensure data integrity. Both encryption and crypto-checksumming are completely transparent to the application, and in both areas the system administrator can chose between several algorithms.

Protecting Data at Rest: Oracle Transparent Data Encryption

Oracle Transparent Data Encryption (TDE) is applied to data in the files which make up the production database (as opposed to backup files, which will be discussed in the next paragraph). As the name indicates, TDE is transparent to the application; no application changes are required. Starting with SAP NetWeaver version 7.20, you can use BRSPACE to set the encryption attribute on a tablespace level. BRSPACE can also be used to manage the wallet, which stores the encryption keys.

Transparent Data Encryption comes in two flavors. The first 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 sensitive data, and encrypt them. Everything else remains unencrypted. The second 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.

Protecting Data at Rest: Oracle Backup Encryption

If you decide to use column encryption, consider to use encryption with your database backup. It is generally much easier to steal backups of the database files than directly from the production database itself. Therefore, another Oracle Advanced Security feature is related to backup encryption.
If you simply backup your database files, only those columns that are encrypted in the production database files are encrypted in the backup

Challenge: In order to read or update data in an Oracle Database that is the data store of an SAP application, the obvious and only choice for legitimate users is this particular application. Attackers, however, who want to bypass SAP's user management and access control, could use either a network sniffing tool to capture data in transit or some kind of file editor to read data at rest, i.e. in a database file copy.

Value Proposition: Oracle Network Encryption (included in Oracle Database EE) protects data in transit. Oracle Transparent Data Encryption and Backup Set Encryption (both part of Oracle Advanced Security) protect data in the production database files as well as their backup copies.

Certification/Support: Oracle Advanced Security is certified for all SAP NetWeaver applications. Implementations are supported by SAP.

Versions: Oracle Database 11g and higher

Implementation: Advanced Security features are activated either via the Oracle Net configuration or using SAP's BR*Tools. Details can be found in SAP Notes 973450 (all versions), 974876 (11g), 2591575 (12c), and 1324684 (all versions)

copies. However, combining Oracle Recovery Manager (Oracle RMAN) and Oracle Advanced Security, whole backup sets (that is, all data) can be encrypted.

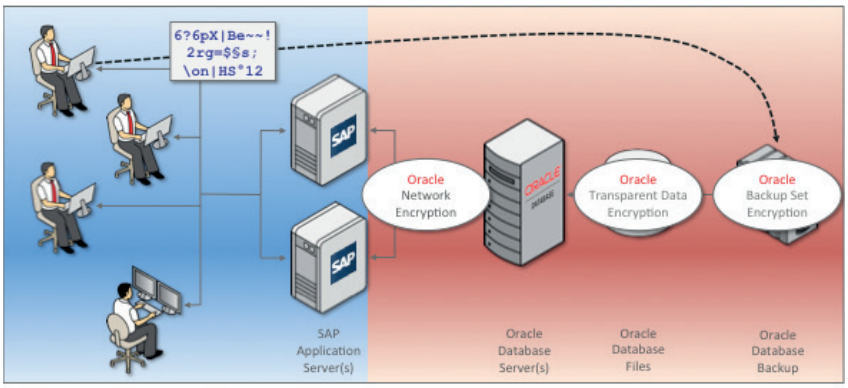

Figure 6.1: Oracle Advanced Security – Encrypt data in transit or at rest

Protecting Data against Unauthorized Access of Privileged Users

Privileged database users – like database administrators – can use DBA tools and directly connect to the database, thus bypassing SAP's security checks (see figure 6.2). Against this kind of threat, encryption does not help. If someone is able to connect successfully using a sufficiently privileged account, and if he or she then sends a query, the Oracle Database will generously deliver the result set to this user. If the requested data is encrypted, Oracle will decrypt it. From an Oracle Database perspective, the request sent by this user seems to be a perfectly valid request.

This can happen, because traditionally, if you were explicitly granted a sufficient number of system privileges, you implicitly received object privileges for all tables as well. For decades people found this acceptable. Recently, however, companies began to ask: Is it really necessary and is it really desirable that a database administrator, who is supposed to manage the database structure, is by default able to read (and even modify) all data in the database?

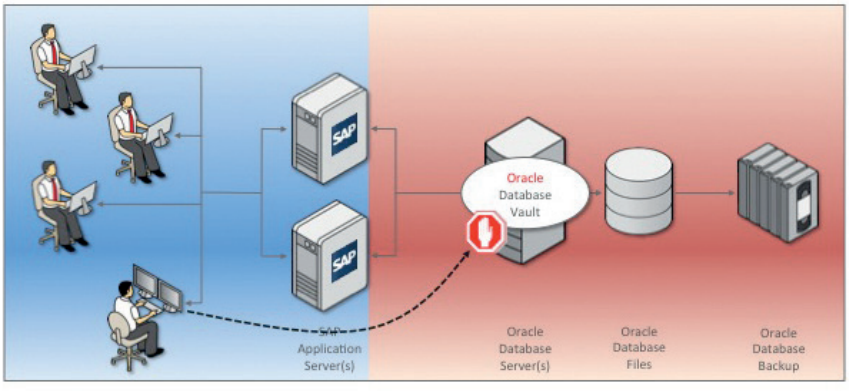

Figure 6.2: Oracle Database Vault – Privileged user access control and analysis

Challenge: Data encryption does not help, if attackers do not use third-party tools, but Oracle tools to bypass the SAP applications. 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.

Value Proposition: 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.

Certification/Support: Oracle Database Vault is certified for all SAP NetWeaver applications. Implementations are supported by SAP.

Versions: Oracle Database 11g and higher

Oracle Database Vault

A solution to this problem requires a new privilege management strategy in the database. This strategy should continue to provide system privileges and object privileges, but it should get rid of implicitly granted object privileges.

This is exactly what Oracle Database Vault does. It replaces the traditional, somewhat clumsy privilege management strategy with a new, more flexible one. It eliminates all implicit grants and instead provides a means to explicitly define access rights as well as the circumstances under which they are effective. This 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.

Implementation: For detailed instructions, see SAP Note 2218115.

Oracle Database Vault for SAP

Oracle Database Vault, as sold by Oracle, is just a toolbox. It is true that it comes with predefined realms and roles, but those are realms for system tables and very general (if fundamental) roles. These predefined components allow Oracle Database Vault to be functional and allow you to use it, but they do not protect your application-specific data. That is because Oracle does not know anything about your applications and your data. Oracle can only give you a toolbox, and it is up to you to determine your security requirements and translate these requirements into an access control policy.

But there is an important difference. As long as customers use home grown applications, it is simply not plausible for Oracle to do more. However, if thousands or tens of thousands of companies use a certain standard application and the security requirements of all these companies are (at least to a certain point) identical, because they are results of the application design, then it makes much more sense for Oracle to analyze the requirements and implement a basic security policy.

Oracle has actually done this, saving customers the time that is needed to implement the boring basics of an application-specific security policy, and protecting these companies from forgetting to implement some of the basics. As of today, Oracle provides a whole family of predefined application-specific Oracle Database Vault policies, and Oracle Database Vault for SAP is a member of this family.

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.

In **traditional auditing**, audit records could be stored in the database audit trail or in files on the operating system. Auditing included operations on privileges, schemas, objects, and statements. To reduce the overhead on the source database system, the audit trail could also be written to operating system files.

Traditional auditing is deprecated starting with Oracle Database 21c.

However, traditional auditing was restricted to the core database functions, and that is why there was a growing number of separate audit trails for additional components, such as Data Pump, RMAN, or Database Vault.

Unified Auditing consolidates all auditing into a single repository and view. This provides a twofold simplification: audit data can now be found in a single location, and all audit data is in a single format. The new policy-based syntax simplifies management of auditing within the database and provides the ability to accelerate auditing based on conditions.

Manageability

INTEGRATED MANAGEMENT TOOLS

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

SAP Migration Tools
For homogeneous or heterogeneous system copies

ORACLE ENTERPRISE MANAGER AND EM PACKS

Enterprise Manager
Comprehensive monitoring and management solution for Oracle Database and Engineered Systems

Diagnostics Pack and Oracle Enterprise Manager
Automatic performance diagnostic and advanced system monitoring

Diagnostics Pack and SAP DBA Cockpit
DBA Cockpit relies on Diagnostics Pack data

Provisioning and Patch Automation
MOPatch utility integrated with Enterprise Manager deployment procedures

ORACLE REAL APPLICATION TESTING

Capture Complete Database Workload
Real world workload instead of artificial workload

Replay Complete Database Workload
Assess impact of system changes before production deployment

Capture Specific SQL Statement
Isolate SQL statements issued by an application

Replay Specific SQL Statement
Predict and prevent SQL performance problems caused by environment changes

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:

- *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.
- *Database migration tools* provided by SAP can be used for homogeneous as well as for heterogeneous system copies.

Enterprise Manager and Enterprise Manager Packs

SAP DBA Cockpit and Enterprise Manager Diagnostics Pack

The Oracle Diagnostic Pack provides automatic performance diagnostic and advanced system monitoring functionality. The Diagnostic Pack includes the following features:

- *Automatic Workload Repository (AWR):* AWR is a built-in repository within every Oracle Database that contains operational statistics about that particular database and other configuration and usage information. At regular intervals, the Oracle Database takes a snapshot of all its performance statistics and workload information and stores it in AWR. AWR forms the foundation for most of the self-management functionality of Oracle Database. It is the source of information that gives the Oracle Database a historical perspective on how it is being used and enables it to make decisions, which are accurate and specifically tailored for the system's environment. Most of the self-managing features of the Oracle Database rely heavily on the information captured in AWR. The data in AWR is also useful for diagnosing all types of performance issues while minimizing administrative overhead.
- *Automatic Database Diagnostic Monitor (ADDM):* ADDM builds upon the data captured in AWR. ADDM makes it possible for the Oracle Database to diagnose its own performance and determine how any identified problems could be resolved. ADDM runs automatically after each AWR statistics capture and makes the performance diagnostic

Challenge: Monitoring and managing a complete IT infrastructure often can be challenging. Administrators may end up with a considerable number of management tools designed for this or that particular purpose.

Value Proposition: Oracle, being a vendor of the whole IT stack from hardware via databases to applications, realized that there is a need for an integrated enterprise management tool. However, in order to avoid unnecessary complexity, Oracle Enterprise Manager Grid/Cloud Control is divided into a base product and several packs dedicated to special features. With or without packs – Oracle Enterprise Manager always provides a complete

data available immediately. ADDM examines data captured in AWR and performs analysis to determine the major issues on a proactive basis, recommends solutions and quantifies expected benefits.

- *Active Session History (ASH):* All active database sessions are automatically sampled once every second and stored in the ASH. The data is captured in a rolling buffer in database memory. The ASH data shows where the database is currently spending its time and highlights any performance bottlenecks. As ASH captures the session state with many performance attributes, the in-memory ASH data can be very effectively used to understand the database workload profile and proactively diagnose any transient performance issue, such as a CPU spike or an I/O storm, that occurs for a very short duration.

Oracle Diagnostics Pack diagnostics and analysis features such as ADDR, ASH, and AWR are tightly integrated with SAP's DBA Cockpit. Therefore DBA Cockpit has a mandatory requirement for Oracle Diagnostics Pack.

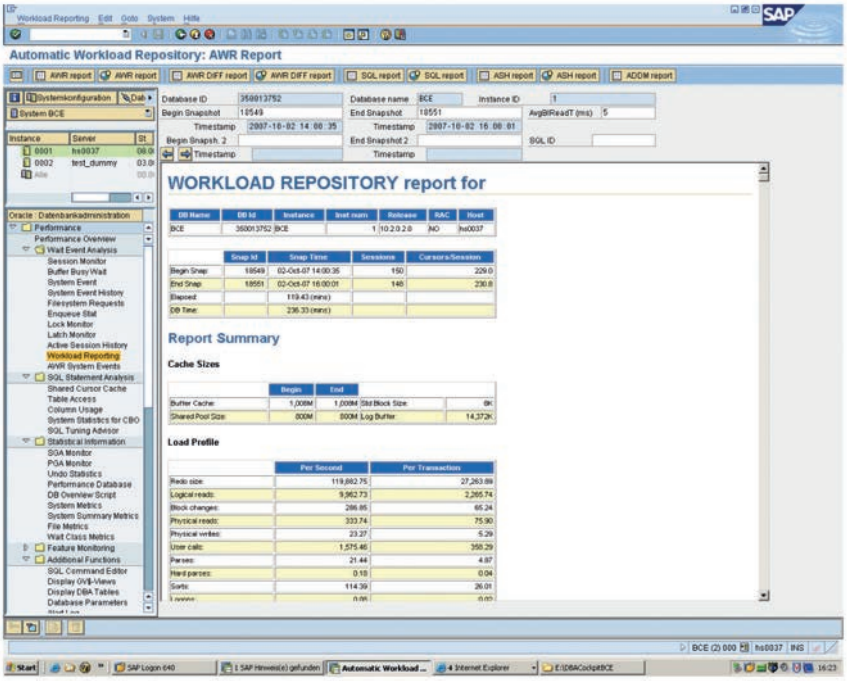

Figure 7.1: Oracle Enterprise Manager Diagnostics Pack and SAP DBA Cockpit

Database Lifecycle Management

The Provisioning and Patch Automation Pack automates the deployment of software, applications, and patches. It makes critical data center operations easy, efficient and scalable, resulting in lower operational risk and cost of ownership. The ability to provision the entire software stack that includes the operating system, middleware and database, supplemented by comprehensive reporting tools, make the Provisioning and Patch Automation Pack an extremely significant entity in overall system management space.

The Provisioning and Patch Automation Pack includes an end-to-end patching solution that works seamlessly across a wide range of products view of the IT landscape.

Certification/Support: Generally speaking, SAP on Oracle customers have a choice. On the one hand there is the BR*Tools family of management tools provided by SAP. These tools include SAP-specific knowledge and are therefore easy to use. On the other hand there are the powerful Oracle Enterprise Manager and its packs. However, in some cases Oracle Enterprise Manager or one of its packs are required – either because there is no BR*Tools support for an Oracle Database feature (example: Database Vault) or because SAP and Oracle functionality are integrated so completely that the former is not possible without the latter (example: SAP DBA Cockpit and Enterprise Manager Diagnostics Pack).

Versions: Oracle Database 11g and higher.

Implementation: For general information about using Oracle Enterprise Manager in SAP environments, see SAP Notes 355770 and 1028068. Special use cases are discussed in additional notes and white papers.

and customer environments. The patching application automates the deployment of Oracle patches for the database and the underlying operating system. The application takes care of pre- and post-patching steps such as shutdown and startup of services, and dictionary changes, if required. It also leverages the flexible Deployment Procedure framework, that lets users add custom steps for specific actions within the patching process.

The Provisioning and Patch Automation Pack also comes with out-of-the-box Deployment Procedures to provision the Oracle Database (both single instance database and RAC), Oracle Clusterware and Oracle Automatic Storage Management from "gold images" following the best practices for maximum availability. The gold images are tested and approved software images and can be patched to any level before deployment.

For SAP environments, the MOPatch utility is integrated with the deployment procedures of Oracle Enterprise Manager to automate the orchestration of patching for Oracle databases.

See white paper "Patching of Oracle Databases in SAP Environments using Oracle Enterprise Manager".

Real Application Testing (RAT)

Oracle Real Application Testing enables you to perform real-world testing. By capturing production database workloads and assessing the impact of system changes before production deployment, it minimizes the risk of instabilities associated with changes. Oracle Real Application Testing comprises two components: Database Replay and SQL Performance Analyzer.

Challenge: Database software patches or upgrades, modified database server configurations, and the implementation of new database features/options can influence performance, availability and security of the database server and the whole system. In particular if the implementation must take into account customer- or application-specific characteristics, administrators will want to know in advance how the new feature or configuration works in the production system.

Value Proposition: The main problem with many test systems is that the workload applied to them is smaller than or different from the workload of the production system, and that a new feature or configuration works well in the test system, but not in the production system. Therefore Oracle Real Application Testing allows customers to capture production database workload and replay it on a test system. Combining these two steps, allows for the true effects of the changes to be understood using a real-life workload before they are implemented in production systems.

Database Replay

Load testing today is generally done using tools that allow testing teams to generate synthetic workloads based on what they expect users to do on a system. These workloads can then be replayed by application virtual users, which simulate the end users by submitting requests to the application. Although widely used, this approach has a number of shortcomings when it comes to testing database level changes:

- Creating the synthetic workload can take a considerable time and requires programming expertise.
- User behavior is not well understood, so many possible workflows are often missed in the synthetic tests.
- Production scale database concurrency is near impossible to simulate with these tools.
- A full application stack is required for testing as these tools simulate end users.

The Database Replay feature included in Oracle Real Application Testing provides DBAs and system administrators with the ability to faithfully, accurately and realistically rerun actual production workloads, including on-line user and batch workloads, in test environments. By capturing the full database workload from production systems, including all concurrency, dependencies and timing, Database Replay enables you to realistically test system changes by essentially recreating production workloads on the test system. This is something that a set of scripts can never duplicate. With Database Replay, DBAs and system administrators can test:

- Database upgrades, patches, parameter and schema changes, etc.
- Configuration changes such as conversion from a single instance to RAC, ASM, etc.
- Hardware and operating system migrations.

SQL Performance Analyzer

Database Replay delivers half of what Oracle calls Real Application Testing; the other half is provided by another tool, SQL Performance Analyzer. The main difference between these two tools is the scope involved: Whereas Database Replay applies to the capture and replay of all activities in a database, SQL Performance Analyzer allows you to capture specific SQL statements and replay them. The latter offers a significant advantage for SQL tuning, because you can tweak the SQL statement issued by an application and assess its impact.

SQL Performance Analyzer (SPA) can predict and prevent SQL execution performance problems caused by environment changes. It provides a granular view of the impact of environment changes on SQL execution plans and statistics by running the SQL statements serially before and after the changes.

Certification/Support: Oracle Real Application Testing is certified for all SAP NetWeaver applications. Implementations are supported by SAP.

Versions: Oracle Database 11g and higher.

Implementation: As this is a database-only feature, no SAP tool support is required. For details, see SAP Note 1426980.

Customer Success

Major Improvement in SAP CRM use with Oracle Database In-Memory

Bosch GmbH accelerates business processes and reduces costs with Oracle Database In-Memory – no changes to existing SAP application required

The Bosch Group is a leading global supplier of technology and services. It employs roughly 375,000 associates worldwide (as of December 31, 2015). The company generated sales of 70.6 billion euros in 2015. Its operations are divided into four business sectors: Mobility Solutions, Industrial Technology, Consumer Goods, and Energy and Building Technology. The Bosch Group comprises Robert Bosch GmbH and its roughly 440 subsidiaries and regional companies in some 60 countries. Including sales and service partners, Bosch's global manufacturing and sales network covers some 150 countries. The basis for the company's future growth is its innovative strength. Bosch employs 55,800 associates in research and development at roughly 118 locations across the globe. The Bosch Group's strategic objective is to deliver innovations for a connected life. Bosch improves quality of life worldwide with products and services that are innovative and spark enthusiasm. In short, Bosch creates technology that is "Invented for life."

Bosch GmbH is known for being a market leader and an innovative company. They often uses the tagline "technology invented for life" which provides insight into their business objectives to improve the quality of life, and to deliver innovation for a "connected life".

For well over a decade, Bosch has relied on Oracle as their preferred database vendor for use with SAP. The Oracle database is a core element and important cornerstone in the Bosch infrastructure. Oracle databases are operated autonomously in combination with the various SAP applications and ideally respond to the varied needs of this successful and diversified group.

The IT department consists of more than 1,500 IT specialists and process experts. It was only logical as an Oracle/SAP customer to fully explore the features and options of the new Oracle Database 12c at an early stage. In mid-2015, the primary focus was on the features and benefits of Oracle Database In-Memory.

"We wanted to find out exactly how Oracle Database In-Memory could help us to implement optimizations in our use of SAP. This applied to both OLAP and OLTP applications. We quickly realized that we stood to benefit in many ways, especially in terms of SAP CRM usage," explains a project manager and database expert from the IT operations organization.

The Bosch IT team discussed solution scenarios and use cases both in-house and with Oracle. It became clear that Oracle Database In-Memory offered significant advantages when used with SAP BW Analytics.
But even more important to Bosch was the fact that Oracle Database In-Memory also enables planned, prioritized optimizations to be achieved in OLTP systems, and specifically with SAP CRM. This was in fact the most pressing problem. The Bosch IT team needed a suitable solution to eliminate sub-optimum performance when using SAP CRM / OLTP at reasonable cost. Oracle Database In-Memory was the right way to go.

The challenge of secondary indexes

In this CRM use case it was mainly secondary indexes, as well as the extra custom-made indexes that had to be regularly created, that proved to be problematic. According to Bosch: “As a result of the intensive and distinctive use of SAP CRM functions by users, the IT team was obliged to continually create extra indexes for certain CRM objects and tables. This was the only way to use the system in line with our individual usage requirements. However, it wasn‘t possible to achieve 100% application coverage in this way, especially in terms of search options and search availabilities in CRM. This is because not all such search options can be supported by additional indexes. Some application functions simply resulted in a timeout.”

Approximately, 80 additional indexes were created and accumulated over time. This resulted in a considerable increased need for maintenance demands and performance related tuning as well as cost-intensive resource use. It was also necessary to set aside a large amount of storage capacity for these indexes, which were also large objects.
The SAP CRM system is used by around 3,500 users throughout the Bosch Group. The CRM system being discussed here is one of several CRM systems Bosch operates, the largest of which has a data volume of over 1TB.

Problems / Challenges:
Bosch was looking for ways to optimize its use of SAP CRM. Specific tables having a large number of secondary indexes were causing particular difficulties. These difficulties resulted in frequent user dissatisfaction, costly administration and maintenance, and the provision of additional IT resources (such as storage).

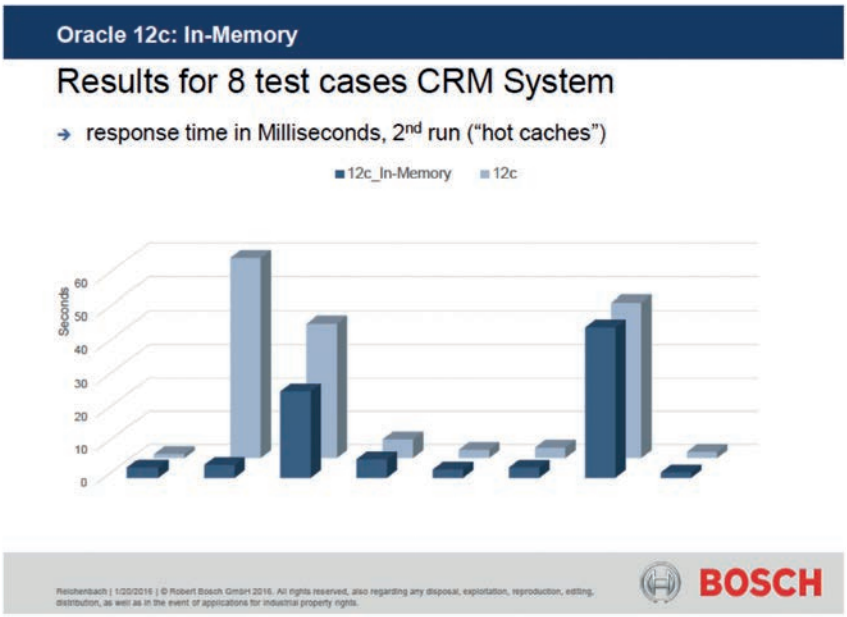

Following a PoC and several tests, it became clear that Oracle Database In-Memory is not only effective as a performance enhancer but it is equally as effective as a means of countering problematic secondary indexes. With Oracle Database In-Memory, many of these indexes can be removed. This prompted Bosch to put Oracle 12c with Oracle Database In-Memory to productive use in conjunction with the SAP CRM system. A detailed analysis was carried out to find out which tables were suitable for In-Memory. The IT team identified them and the most important ones were able to be replaced. One major advantage is that the benefits of Or-

Solution:
With Oracle Database In-Memory, it is possible to replace additional custom-made indexes. Oracle Database In-Memory is used for the tables that made these indexes necessary.

acle Database In-Memory technology can be gained without any changes to SAP applications, with the continued use of existing hardware (virtualized and Unix-based servers), and without data migration. It is usually only necessary to provide In-Memory servers / RAM hardware resources (virtualized or non-virtualized) for tables that were to be involved with In-Memory processing.

High satisfaction with Oracle Database In-Memory

„The improvements achieved in the SAP CRM system through Oracle Database In-Memory have turned out exactly as we planned. User satisfaction has risen considerably and IT-supported business processes are now running smoothly without delays and can be fully utilized as standard. The amount of IT administration has also been noticeably reduced thanks to the substantial reduction in index creation and maintenance, which in turn results in cost savings.

In the use case described, Oracle Database In-Memory eliminated the need for many secondary and custom-made indexes of large tables while accelerating OLAP use. Inserts, Updates and Delete statements processed normally. No negative impact on the system performance was observed. Additional tuning can be accomplished easily by increasing the cache for the In-Memory Store and assigning specific tables to occupy the space. The Oracle In-Memory Advisor, a component supplied with Oracle Database, also makes it possible to identify tables suitable for use with In-Memory.

Benefits:
Now users can make full and effective use of SAP CRM

- Reduced costs for system / database administration
- In-Memory is only activated for selected tables
- Saving in terms of storage resources for secondary indexes that are no longer needed
- Lower number of indexes means accelerated business processes
- Faster SAP CRM response times and higher system performance
- Use of Oracle Database In-Memory without changes to SAP applications
- Continued use of existing hardware; only In-Memory / RAM must be provided

SAP BI with Oracle Database In-Memory at DB Masters

In-Memory database technology has been around since the 1990's. Numerous database vendors have come and gone and filled that niche for years. In 2011 SAP came up with the HANA In-Memory concept. The initial idea was to efficiently handle analytical queries through columnar stored data structures. Planned as a BWA replacement, the technology matched well.

Later, SAP expanded the concept to include OLTP. Technically, they chose to mix self-developed and acquired technologies such as P*Time, TREX Search Engine, MaxDB (livecache), and Sybase ASE for this purpose. Customers wanting to exploit the In-Memory technology from HANA require special hardware to run the system and a migration/conversion from the technology they knew today to the HANA flavor of the day. Nearly at the same time, Oracle fully integrated the In-Memory technology into the Oracle 12c database.

Oracle Database In-Memory promises high performance of complex analytical processing, near transparent implementation, and requires no special hardware or data migration. When done properly, customers will continue to benefit from the performance, scalability, availability, and reliability available today with the added high performance analytical capability.

This article explores the pre-certification testing and findings of one customer's experience with Oracle Database In-Memory. Oracle Database In-Memory has been certified for use with SAP since June of 2015.
We at DB Masters in Austria are focused on solving customer business challenges by proper implementation of the Oracle Database and related technologies. We have experience with many large ERP systems and understand their special needs. However, our focus is Oracle centric - independent of the applications running on it. We are in business since 2000, servicing customers across Europe. Our customer list is long.
When Oracle 12.1.0.2 In-Memory had been initially released by Oracle, one of our customers expressed interest in this technology. The company approved to a proof of concept to evaluate the performance impact for their SAP BW application.

We took the "top" (longest running) SQL statement from their production SAP BI and captured measurements from the following scenarios:

- The "current" query runtimes from the production environment as baseline – which is not entirely accurate as there are a lot of other activities in the database.
- 11gR2 with both the original buffer cache size and 256GB cache. – Statements were executed several times to warm up the cache, so we were able to eliminate the I/O time from the execution. This allowed us to measure performance differences between Oracle 11gR2 and Oracle 12c results.
- 12cR1 with both the original buffer cache size and 256GB cache. – Statements were also executed several times to warm up the cache so we were able to eliminate the I/O time from the execution. This allowed us

About the Author:
Christian Pfundtner, CEO of DB Masters, is working with Oracle Databases since 1992 (Oracle 6). One of the first 4 OCMs in Europe. Our focus is on the Oracle Database – but unlike many of our competitors we take care of all related layers – from the disk (storage) up to the user. We offer all kinds of evaluations. Business and current database requirements, current licensing state, recommendations on best practices for implementing Oracle to address your needs and to maximize your ROI, performance tuning and troubleshooting.

to measure performance differences between Oracle 11gR2 and Oracle 12c results. The runtime difference between Oracle 11g and 12c were negligible. We can assume there are no performance differences for the test statements.

- Enable Oracle Database In-Memory and put all related objects into the In-Memory Column store. Several configuration settings were tried – like different In-Memory compression options – but at the end all results were done with FOR QUERY LOW compression.

Here are the results over all queries – keep in mind that this is an older SAP BI Version, not utilizing the optimized flat cube model for Oracle Database In-Memory technology:

Type of Test	Result
Runtime on production system	Always between 2500 and 2750 seconds depending on the current system load
11g / 12c Test system original buffer cache	About 1455 seconds
11g / 12c Test system 256GB buffer cache	About 382 seconds
Using Oracle Database In-Memory with QUERY LOW Compression	About 118 seconds

The numbers may not be as impressive as when Larry announced the Oracle Database In-Memory, but depending on the baseline running up to 25 times faster does not look so bad.

There still remains the question often coming: “Why do you not put the whole database into buffer cache?” This would speed up the system by a factor of 7-8.

The answer is easy: most customers will not be willing to pay for servers having the required amount of memory. If your SAP BI database is 5, 10 or more TB in size, you would still need up to 50% of the database size for buffer cache. Today, there are servers having up to 32TB and even more memory, but they are very expensive. Current Intel based servers with 2 sockets are able to address up to 3TB memory – typically affordable servers are populated with up to 1.5TB memory.

The advantage of Oracle 12c Database In-Memory is that only table data must fit into the SGA. This typically ranges between 1/3 to 1/2 of the database size, depending on the type of system. The fact that In-Memory table data is stored and accessed in a compressed form means that the footprint will be reduced further. Our tests indicate that most SAP BI tables could be compressed by a factor between 2 to 4x’s of the original table size. In other words, the size of the In-Memory table will be 1/2 to 1/4 of the original table size.

A current Intel based server with 2 sockets and 1.5TB memory should be able to put a 10TB SAP BI database completely into the memory! Additionally, you can get rid of indexes only used to speed up lookups – es-

pecially bitmap indexes needed for query performance. Fewer indexes to maintain will speed up data loading processes in a dramatic way while reducing disk space and backup sizes.

Another advantage of Oracle Database In-Memory is the short amount of time it takes to implement it. Believe it or not, this customer POC was done within 5 hours – including creating a clone of the database on the test system and running all of our tests! You may compare this to any SAP Feature implementation project you like to.

Deactivating the Oracle Database In-Memory is also quite easy. One simple command line directive dynamically enables and disables the feature. This offers a quick and efficient fall-back if any show stopping issues should arise.

Last but not least – how is this effecting HA/DR? With Oracle database In-Memory, you may continue to still utilize all High Availability functionalities like Oracle RAC, Oracle Data Guard, and fail-over solutions from Oracle or other vendors. No changes. Recoverability remains also untouched. The feature is fully integrated with all of the Oracle HA and DR capabilities.

In summary, experience with this new Oracle feature was very positive. It exceeded our expectations with respect to the ease of implementation and the performance benefits. It is a certified option and nearly transparent to implement. Customers now have a viable option to explore this technology at a reasonable cost with no business disruption.

DB Masters GmbH
Stammersdorfer Str. 463
2201 Gerasdorf, Austria

www.dbmasters.at

Cost Savings and many other Improvements with Oracle Database and Oracle Database In-Memory

kommunale
informationsverarbeitung
baden-franken

KIVBF is an IT systems vendor and a complete solutions provider for cities, municipalities and districts.
Being efficient, secure and innovative from a commercial standpoint, the KIVBF portfolio covers the complete range of municipal data management for finance, HR, regulation and reporting with contemporary specialized IT solutions and services.

Extensive ranges of consulting and training services, as well as IT-services and ASP solutions generate sustainable added customer value. We believe that specialized IT procedures and lively dialog with decision-makers and users are critical for meeting the short and long-term business requirements.

As a complete solutions provider, we advise and support our customers from the conceptual design process through implementation, and offer ongoing application and user support. We optimize the workflow through automation and optimized connections to existing procedures or external agencies.

Our customers receive a complete and comprehensive service that satisfies the highest technical standards for the ongoing optimization of their business processes. We analyze trends and respond to current and future legislative requirements and amendments at an early stage. In doing so, we offer our customers the solutions they need in a timely manner: through working together as partners!

Tool for optimizations using Oracle Database for SAP

In 2015, one of our customers requested our assistance in a campaign to optimize their IT infrastructure with the following goals in mind:

- Migration of the predominantly used MaxDBs to Unicode and replacement of the mainframe and the associated DB2 databases of a larger SAP IS-U.
- Replacement of the existing SAP BW infrastructure (including BWA) with a more powerful and more modern solution.
- Revision of the DB infrastructure. It had been fore-casted that there would be a 30% growth in data as a result of migrating the MAXDB systems to Unicode.

KIVBF analyzed the current and future requirements along with the total cost of ownership. The best value for the money was Oracle. After performing some fundamental tests, KIVBF decided to replace the mainframe systems with Oracle Database 12c – which was urgently needed. According to KIVBF, the Oracle DBMS satisfied their requirements with respect to high availability, disaster recovery, security, performance and stability. Oracle was the best option for a powerful, stable, secure, cost-efficient and future-proof database platform for SAP NetWeaver based applications.

As Uwe Bersch from KIVBF's specialist "Databases" division explains, the replacement of the DB2/zOS mainframe databases with Oracle x64 Windows was a smooth process from both a technical and a time-related standpoint. By using the latest Oracle technologies and "Advanced Compression" in particular, performance was given a considerable boost, the data volume was halved, and the backups were three times smaller. These facts can be presented well in economic terms too!

Extra added value following migration through use of Oracle In-Memory

Due to its positive experiences and economic situation, KIVBF decided to put Oracle 12c in place as a strategic SAP database. KIVBF also dealt a great deal with the Oracle 12c In-Memory Option, which has been certified by SAP for all SAP NetWeaver-based applications since 2015. The BW NetWeaver (BWA) which had been used up until then was planned to be phased out with the switchover.

SAP HANA was also considered as an alternative to Oracle. The question as to whether the Oracle In-Memory Option or SAP HANA was suitable had certainly been evaluated. "In the end, it was the cost factors in particular that tipped the scales and led us to choose Oracle In-Memory for our SAP BW operations," reported KIVBF.

To be specific, KIVBF would have had to buy new specialized hardware for SAP HANA, and it would also have had to pay extra HANA license costs too. Another technical advantage that Oracle In-Memory has over SAP HANA is that Oracle In-Memory enables the user to upload only selected, performance-relevant tables. The Oracle DBMS can therefore get by with a more streamlined working memory.

The migration of the SAP BW NetWeaver 7.31 database to Oracle 12c, using Oracle Database In-Memory, and the replacement of the SAP BWA took just eight weeks. The solution went live in December 2016, with the DB2 and MaxDB migration processes taking place almost in parallel. A toolkit developed by Oracle was used in conjunction with the Oracle 12c In-Memory Option to optimize the BW data processing. The toolkit essentially creates Materialized Views*, which present a mock-up of the InfoCube tables and can thus be loaded to the "In-Memory" in a "de-normalized" state. Oracle can now respond to SAP Netviewer inquiries flexibly by means of forwarding the requests or parts of them to the Materialized-Views (RE-WRITE). "Materialized Views" & "Query Rewrite" are a highly flexible technology that has been provided by Oracle for 15 years to optimize complex queries. In parts, the technology is similar to the Flat Cubes introduced with NetWeaver 7.4, which guarantee the In-Memory technology's huge performance gains.

∑ SAP MaxDB with BWA in seconds	2899
∑ Oracle (Materialized V+IM) in seconds	335
Speed ratio Oracle (IM) : MaxDB (BWA)	8.65

Table 1:
Speed comparison: SAP MaxDB with BWA and Oracle 12c In-Memory

* http://www.oracle.com/us/solutions/sap/sap-database/database-in-memory/index.html
Please see SAP Note 2351252

Oracle Customer:
Kommunale Informationsverarbeitung Baden-Franken (KIVBF)

www.kivbf.de

Industry: Service provider (public service)

Turnover: Over 1 billion

Workforce: 610

SAP MaxDB (uncompressed)	2.48
SAP MaxDB (uncompressed) with Unicode	3.2
Oracle 12c (compressed) with Unicode	0.94
Savings factor	3.4

Table 2:
Comparison of BW database volumes: SAP MaxDB and Oracle 12c in TB

Oracle In-Memory makes SAP BW reports up to 36 times faster

According to the Kommunale Informationsverarbeitung Baden-Franken (KIVBF), the move towards Oracle 12c including the Oracle In-Memory Option was absolutely worthwhile: "We achieved and even exceeded all the goals we set," explains Uwe Bersch, of the Technical Solutions and Database Service division.

The municipal IT service provider is completely satisfied, particularly with the use of Oracle In-Memory. Firstly, it was able to save on license costs – particularly due to the replacement of the SAP BW Accelerator – and substantially reduce the BW data volume (in the backup too) – which ultimately also translated into cost savings. While at the same time they were able to significantly boost performance using the Oracle In-Memory Option.

The BW report response times have become three times better on average in evaluations with runtimes of 10 to 20 seconds with Oracle In-Memory when compared with the SAP MaxDB including BWA used previously. Complex evaluations with runtimes ranging between 3 and 12 minutes have even been made 11 to 36 times faster. The BW Oracle Database now has a (compressed) data volume of less than 1 TB; previously, it was 2.48 TB (with Unicode and MaxDB, the data volume would have increased to 3.2 TB). The backup volume, meanwhile, is just 330 GB – thanks to Oracle. The BW data loading times have been significantly shortened with Oracle too. KIVBF currently has four ERP systems from which the BW system gets its data for reports / evaluations.

As Brigitte Landwehr, the BW system's specialist manager, explains, her colleagues in new customer projects have told her "that they were amazed at how quick the BW response times are with Oracle In-Memory". She went on to add that some reports for In-Memory processing had not been used because they "achieved a performance in the low single digits using Oracle 12c alone, even without In-Memory".

KIVBF employees certified by SAP are presently (in the middle of March 2017) switching the remaining SAP systems over from SAP MaxDB to Oracle 12c. Yet another priority item on KIVBF's optimization agenda is switching SAP NetWeaver 7.31 to either version 7.4 or version 7.5. In this regard, the support of Flat Cubes in BW In-Memory processing, which supports Oracle 12c with OIM from the version SAP NetWeaver 7.40, is significant. This solution not only optimizes performance for analyses, but considerably minimizes loading times too.

Oracle Products & Services:

- Oracle Database 12c (12.1.0.2), with Oracle Advanced Compression and Oracle In-Memory Option for SAP NetWeaver BW
- Oracle RMAN (Compression) for backup
- Oracle Advanced Customer Service (part of Oracle Customer Success Services, as of July 2023)

Key benefits:

- Cost optimizations
- minimized database footprint
- much better performance
- future-proof database technology

SAP:

- SAP Industry Solution IS-U/ ERP ECC, NetWeaver BW, and others

Infrastructure:

- Cisco server (virtualized VMware)
- NetApp storage systems
- OS Windows Server 2012 V2

Tremendous Improvement in Performance in SAP BW with Oracle Database In-Memory and “Flat Cubes with Oracle”

With an annual turnover of approximately 820 million and over 7,700 employees, Villeroy & Boch is one of the world‘s leading premium brands for ceramic products. The family business, which was founded in 1748 and is headquartered in Mettlach/Germany, stands for innovation, tradition and exceptional style. As a renowned lifestyle brand, Villeroy & Boch offers products from the sectors Bathroom and Wellness and Tableware, and is active in 125 countries.

Both business areas use an SAP ERP and an SAP BW system as well as other SAP solutions. Villeroy & Boch has also been using Oracle’s proven, reliable database to great advantage for around 20 years. Over the years the use of BW/BI has steadily increased. The company has approximately 2,500 SAP users worldwide, including some 1,700 in the SAP BW environment. Certified SAP hosting partner Rödl IT Operation GmbH (with offices in Mettlach, Selb, and Frankfurt) operates and manages a range of systems including all SAP systems on behalf of Villeroy & Boch AG.

Putting Oracle Database In-Memory to the test

Shortly after the Oracle Database In-Memory technology was officially certified by SAP in the middle of 2015, Villeroy & Boch joined forces with Rödl & Partner, its hosting partner, to put Oracle Database In-Memory to the test. This new database 12c technology for its SAP BW was extensively tested on a sandbox system on the exact same hardware as used with the productive system. During this process, the fact tables (E and F tables) of the biggest InfoCube (COPA) – containing approximately 94 million complete datasets – were loaded in the Oracle database’s additional, column-oriented memory area “Column Store”. The tests mainly involved directly comparing pure read times for selected queries with the productive system. The database system needed to be upgraded from Oracle 11g to Oracle 12c to use Database In-Memory. The database’s main storage was In-Memory. The database’s main storage was expanded by 96 GB. This new memory is only used in the database for column-oriented In-Memory tables. The buffer sizes in the System Global Area (SGA), such as those for the data block buffer cache and shared pool remained unchanged. The entire test took just 1.5 days. Only three Oracle database parameters are needed to use Database In-Memory.

- inmemory_max_populate_servers=4
  (limits the number of database background processes for loading the In-Memory tables, default = 4)
- inmemory_size=96 GByte
  (defines the size of the In-Memory Store)
- inmemory_clause_default = „PRIORITY HIGH“

The SAP LISTCUBE transaction, to name but one example, was used when carrying out the tests. The database response times were calculated with SQL traces using SAP transaction ST05 (performance analysis). All the database tables were compressed on the productive system using Oracle Advanced Compression technology. The indexes, meanwhile, were compressed with Index Compression, which has been usable in SAP since Oracle 10g. Both compression operations therefore took place using the test system too.

Expectations exceeded

Comparisons with the productive system were made on identical hardware with various test scenarios on the SAP NetWeaver BW system.

- In the first test, 1 million datasets were selected without database aggregation several times over. While 24 seconds were measured for this query in the productive system, a response time of 160 milliseconds was recorded on the identical test system with Oracle In-Memory.
- These 1 million datasets were aggregated in the second test too. Here, the execution time was shortened from 44 seconds to 11 seconds.
- In the third test, the entire data volume of the loaded tables (94 million datasets) was read and aggregated. While the execution did not give a result in the productive system due to a time-out, it was completed in 22 seconds on the test system.

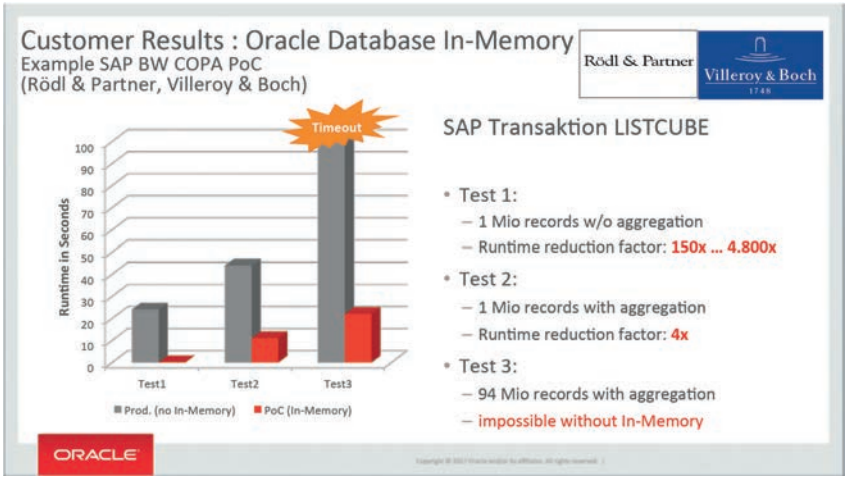

Figure 1: Performance of Oracle Database In-Memory

In further tests, two existing COPA queries were used and compared without any adaptations.
Here, the execution times were greatly accelerated, going from 50 to 1.5 seconds and from 72 to 2.5 seconds.
The now-shorter execution times are particularly valuable for Villeroy & Boch, as is the modest amount of implementation work required. Use of Oracle Database In-Memory now enables selection and aggregation of extremely large amounts of data in far shorter times.

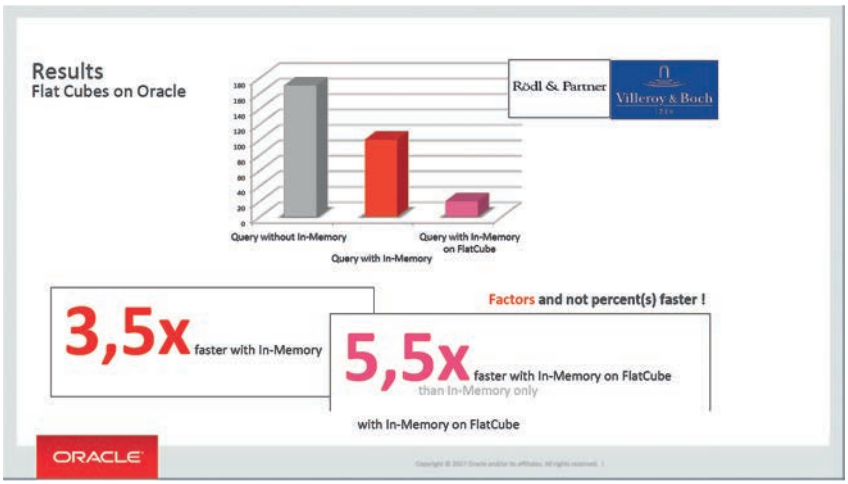

Figure 2: Performance of Oracle Database In-Memory for COPA queries

The moderate amount of implementation work is a key argument in favor of using Oracle Database In-Memory. In a nutshell, only the following was necessary:

- Use of existing hardware, operating system, and virtualization technology
- Database upgrade from Oracle 11g to Oracle 12c
- Activation of Database In-Memory using a few database parameters
- Activation of tables for In-Memory
- No table or database migration
- No adaptations to the SAP application code

Low resource requirement, many benefits

Villeroy & Boch was expecting performance to increase substantially. This much has certainly been confirmed – and impressively so. At the same time, it became clear just how much potential the Oracle Database In-Memory technology harbors.

The swift implementation and performance of the tests both make a completely convincing argument. The fact that there was no need to buy any brand-new hardware given the moderate additional memory costs was yet another positive aspect in favor of using the technology. Use of Oracle Database In-Memory is activated only for tables that are useful for this purpose. It is not appropriate to equip all the tables in a database with In-Memory technology. Doing so makes tables smaller, or even creates tables on which a pure single set query is performed. Column-oriented storage with Oracle In-Memory is not necessary for this.

The technology can be used immediately in the existing infrastructure and virtualized environment. No special adaptations or modifications to the existing database or within the SAP application are necessary. The administrative work involved is also very moderate. The minimum risk involved in upgrading the Oracle database is well known due to many years of experience in dealing with the Oracle database.

- Oracle Database In-Memory results in a significantly improved report execution time
- Improvements several times over
- Moderate implementation work
- Implementation on existing hardware
- No modifications necessary in the application
- No data migration required
- Unmodified use of all Oracle database functionalities

Further developments in the SAP NetWeaver BW now available with Flat Cubes and Oracle too

When it launched SAP HANA, SAP developed the HANA-optimized InfoCubes, which was only possible due to the use of column-oriented In-Memory technology. The "SnowFlake" data model of the existing InfoCubes made up of fact and dimension tables was dissolved and migrated to a flat table.

The same technology has been available to all SAP NetWeaver BW customers since the middle of 2016 along with the Oracle database and the In-Memory option under the keyword "Flat Cube on SAP BW on Oracle". Existing InfoCubes are converted as needed using the SAP standard program RSANAORA. By means of re-partitioning, the InfoCube (E and F fact tables and all the associated dimension tables) are being created as a single flat table in a partitioned format.

See SAP Note 2335159

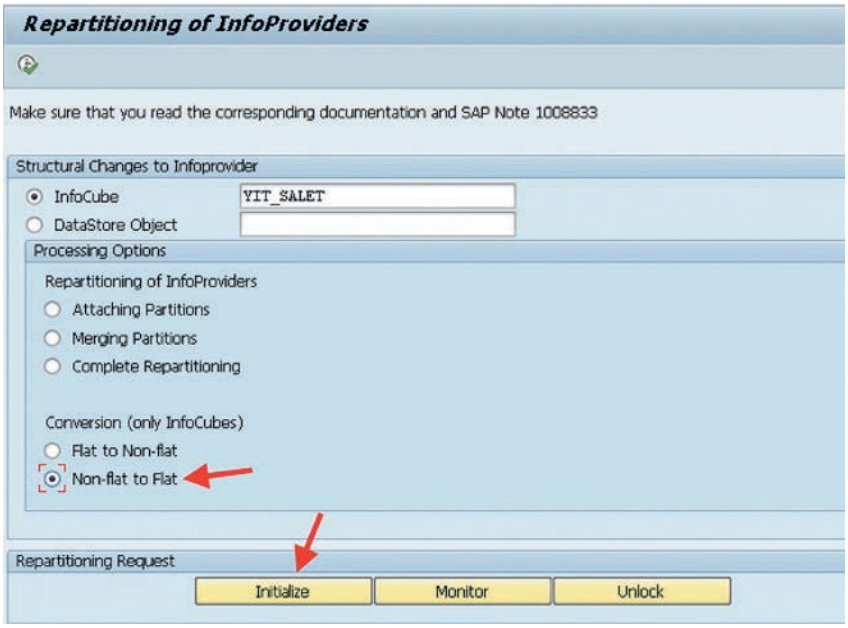

Figure 3: Re-partitioning and activation of the Flat Cube

Following the productive launch of Oracle In-Memory technology in the middle of 2016, Villeroy & Boch converted some InfoCubes to Flat Cubes in SAP NetWeaver BW during another test conducted at the end of 2016 / start of 2017.

So, for instance, an InfoCube with 110 million datasets and measuring 13 GB in size (7 GB of which was index) needed one hour and 42 minutes for the conversion process. Once the conversion was complete, the Flat Cube was 7.2 GB in size and had a 0.3 GB index. In addition to the improvement in response time, which again was to be expected, there was also a reduction in database size within storage.

Further significant increase in performance thanks to Flat Cubes in Oracle

Following the conversion, the execution times of 100 seconds on the previous In-Memory solution could be substantially improved yet again, being shortened to 20 seconds with a Flat Cube and In-Memory. What's more, the benefits also include approximately 30% faster loading, eliminated aggregates and simplified BW modeling – and all of this is achieved through using a Flat Cube.

Oracle recommends using the SAP NetWeaver version 7.40 with SP16 for conversion to Flat Cubes. In 2016, it was still impossible from an SAP standpoint to convert transactional (planable) InfoCubes or semantically partitioned objects (SPOs).

Support for Semantically Partitioned Objects has been added in 2018.

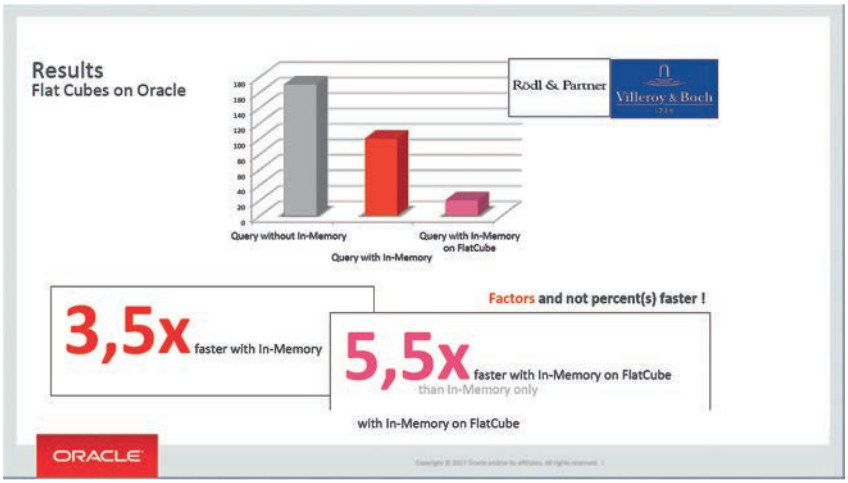

Figure 4: A comparison of execution times

Summary of the tests

Using Oracle Database In-Memory, there is currently no need for Villeroy & Boch to replace its Oracle database for the current SAP NetWeaver systems, which has been reliable for many years, with another database system. Villeroy & Boch is successfully using the Oracle Database In-Memory technology for the productive SAP NetWeaver BW as well as SAP CRM systems.

With a reasonable amount of work and acceptable costs for the investments, the following benefits very quickly became a reality thanks to the Oracle Database In-Memory technology:

- several significant increases in the speed of query times for reports
- far more effective analysis options in reporting

The Oracle database is thus once again proving its position as a long-term stability factor for databases in the SAP environment and delivers innovations that can be implemented very quickly and with very little effort.

"With our tests based on Oracle 12c and Oracle Database In-Memory, the PoC produced amazing results in terms of improved performance. From a BI point of view, I was also especially pleased with the fact that once SAP applications have been created, they can be used with Oracle Database In-Memory without any modifications. It's also easy to choose which SAP BW tables are processed with In-Memory and which are not."

Harald Wolf
Senior Consultant for BI
Villroy & Boch

Oracle-related SAP Notes

Oracle Database-related SAP Notes

Note No.	Note Title	DB Version
DB : General : Licensing Information		
581312	Oracle Database Licensing Restrictions	12c-19c
740897	Oracle License Scope & Required Oracle Options	12c-19c
1028068	Required Oracle Options for DBA Cockpit	12c-19c
DB : General Version Support		
1174136	Oracle: End of Support Dates	12c-19c
2606828	Oracle Database Roadmap for SAP NetWeaver	12c-19c
DB : General : Oracle Support for HANA-Related SAP Application Optimizations		
1835008	Activate Database Performance Optimizations for SAP ERP	12c-19c
1892354	SAP Strategy for Cluster and Pool Tables	12c-19c
1951491	Minimal DB Platform Requirements for SAP NetWeaver 7.4 SP08	12c-19c
DB : Features : Overview		
105047	Support for Oracle Functions in the SAP Environment	12c-19c
2799900	Central Technical Note for Oracle Database 19c	19c
2817074	Oracle Database 19c: Integration in SAP Environment	19c
DB : Features : File Systems		
155013	Using Oracle ASM with SAP NetWeaver-based Products	12c-19c
1416773	Oracle Direct NFS	12c-19c
1627541	BR*Tools Support for Oracle ASM and Exadata	12c-19c
1738053	SAPinst for Oracle ASM Installation	12c-19c
1825328	Oracle ASM and SAP NetWeaver Master Data Management (MDM)	12c-19c
2163827	LVM support for Oracle RAC and/or ASM	12c-19c
DB : Features : Performance		
741478	FAQ: Materialized Views	12c-19c
DB : Features : Data Storage & Compression		
1109743	Use of Index Key Compression for Oracle Databases	12c-19c
1426979	SecureFiles – The New Way to Store LOB Data	12c-19c
1464156	BRSPACE 7.20	12c-19c
1583303	Deferred Segment Creation	12c-19c
1842044	Dataload Without Dropping Bitmap Indexes	12c-19c
1856270	Performance Improvements for Tables with Single Unique Index	12c-19c
DB : Option : Advanced Compression		
1289494	FAQ Oracle Compression	12c-19c
1431296	LOB Conversion & Table Compression with BRSPACE 7.20	12c-19c
1436352	Advanced Compression for SAP Systems	12c-19c
2138262	Oracle Database 12c Advanced Compression for SAP Systems	12c-19c
2157904	Oracle 12c: Conversion of Compressed Tables	12c-19c
2166836	Oracle 12c: Problem during SAP Upgrade with Compressed Tables	12c-19c
2254836	BR*Tools Support for Oracle ADO/ILM	12c-19c
2254866	Using Oracle Database 12c Automatic Data Optimization	12c-19c
2255992	R3load and R3szchk: New Oracle Feature for Database ILM Policy	12c-19c

2258061	Enhancements for ADO/ILM for Table Conversions or System Copy	12c-19c
2384534	LOB Conversion & Table Compression with BRSPACE 7.40	12c-19c
DB : Option : In-Memory		
2137032	DBA Cockpit: Monitor for In-Memory Feature	12c-19c
2178980	Using Oracle Database In-Memory with SAP NetWeaver-based Products	12c-19c
2189163	Oracle Database In-Memory Advisor for SAP	12c-19c
2335159	Flat Cubes for SAP BW on Oracle DB	12c-19c
2351252	Oracle Database 12c In-Memory Toolkit for SAP BW	12c-19c
2711358	Conversion of Semantically Partitioned Objects to a Flat Format	12c-19c
2816467	Support for FEMS Pushdown on Oracle	12c-19c
2948714	BW Migration to BW on Oracle by Respecting Flat Cube	12c-19c
DB : Option : Multitenant		
2333995	BR*Tools Support for Oracle Multitenant Database	12c-19c
2335850	Transformation of Existing Standalone Database into Pluggable Database	12c-19c
2336881	Using Oracle Multitenant with SAP NetWeaver-based Products	12c-19c
3144616	Using Oracle Multitenant with Real Application Clusters (Pilot Program)	19c
DB : Option : Partitioning		
742243	General Table Partitioning	12c-19c
859841	Deinstallation of Oracle Partitioning Option (SAP BW)	12c-19c
DB : Option : Real Application Clusters (RAC)		
196927	Protecting SAP Instances through Oracle Clusterware (SAPCTL)	12c-19c
527843	Oracle RAC Support in the SAP Environment	12c-19c
1431241	RAC with TAF	12c-19c
1977463	Using Shared or Local Oracle Homes with RAC	12c-19c
DB : Option : Real Application Testing (RAT)		
1426980	The Real Application Testing Option	12c-19c
DB : Features : Security		
1157499	OPS$ Connect and Security Aspects	12c-19c
1622837	Secure Connection of AS ABAP to Oracle via SSFS	12c-19c
1623922	Connect to Oracle Database	12c-19c
1710997	Using Personalized Database Administrator Accounts	12c-19c
1755636	Database Administrators Segregation	12c-19c
1764043	Support for Secure Storage in BR*Tools	12c-19c
1868094	Overview: Oracle Security SAP Notes	12c-19c
3130747	Apache Log4j vulnerability	12c-19c
DB : Option : Advanced Security		
973450	Network Encryption	12c-19c
1324684	Backup Encryption	12c-19c
2591575	Transparent Data Encryption	12c-19c
DB : Option : Database Vault		
1355140	Using Oracle Database Vault in an SAP Environment	12c-19c
1502374	Database Vault for SAP Policy Scripts	12c-19c
1716939	SAP System Copies with Database Vault	12c-19c
2218115	Oracle Database Vault 12c	12c-19c
Deployment : Software : Installation & Upgrade		

819829	Oracle Instant Client Installation and Configuration on Unix	12c-19c
1915299	Troubleshooting Software Installation	12c-19c
1915317	Migrating to Software Owner ‚oracle'	12c-19c
1915323	OS User Concept for Oracle Database 12c and higher	12c-19c
2184253	Maintaining SAP and Non-SAP Databases on one Oracle Installation	12c-19c
2800001	Upgrading to Oracle Database 19c (DBUA)	19c
2800021	Upgrading to Oracle Grid Infrastructure 19c	19c
Deployment : Software : Patches		
1027012	MOPatch – Install Multiple Oracle Patches in One Run	12c-19c
1696869	Patching of Oracle Homes with Minimal Downtime	12c-19c
2011199	SAP Bundle Patches – Structure Changes Starting June 2014	12c-19c
2574651	New Strategy for Oracle Database Patching	12c-19c
2799920	Database: Patches for 19c	19c
2799930	Grid Infrastructure: Patches for 19c	19c
Deployment : Instance Configuration		
1171650	Automated Oracle DB Parameter Check	12c-19c
1431751	Quick Reference for ADRCI and ADR	12c-19c
1589924	SAP Workload Management via Oracle Resource Manager	12c-19c
2470718	Database Parameters 12.2 / 18c / 19c	12c-19c
Deployment : Admin Tools : SAP BR*Tools		
12741	Current Versions of BR*Tools	12c-19c
1428529	Corrections in BR*Tools Version 7.20	12c-19c
1598594	BR*Tools Configuration for Oracle Installation under „oracle" User	12c-19c
1838850	Corrections in BR*Tools Version 7.40	12c-19c
1844380	Minor Functional Enhancements in BR*Tools (5)	12c-19c
1882182	Functional Enhancements in BRARCHIVE 7.40	12c-19c
2086966	Minor Functional Enhancements in BR*Tools (6)	12c-19c
Deployment : Admin Tools : Oracle Enterprise Manager		
355770	Oracle Enterprise Manager Database Control	12c-19c
386413	Oracle Enterprise Manager (OEM) in the SAP Environment	12c-19c
Deployment : Data		
646681	Reorganizing Tables with BRSPACE	12c-19c
937492	Oracle Flashback FAQ	12c-19c
966117	Oracle Flashback Database Technology	12c-19c
Database : Services		
1508271	Oracle Lifecycle Migration Service (formerly O2O/Triple-O)	12c-19c

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

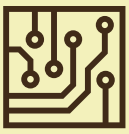

Best Integration of Hardware & Software

... 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

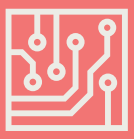

Best Integration of Hardware & Software

ORACLE

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores,
CA 94065 U.S.A.
Internet: www.oracle.com/sap
https://www.sap.com/community/topic/oracle.html