

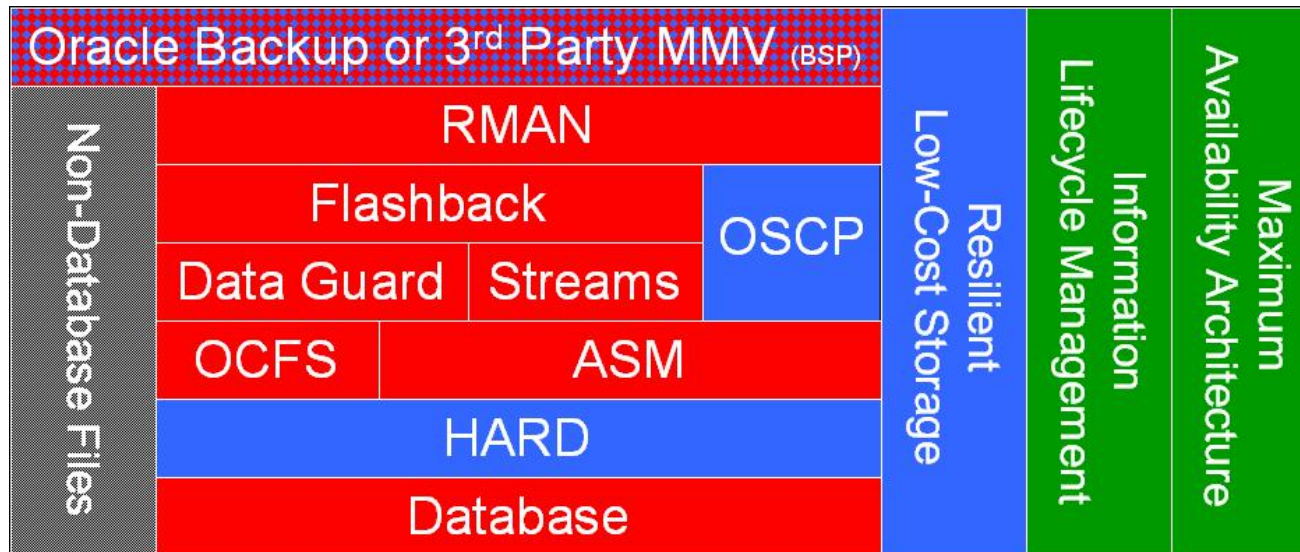
## **An Oracle Technology Brief: The Oracle Database and Storage**



The Oracle instance lives in main memory, but the Oracle database resides on disk storage. One of the founding principles of database systems is that transactions are durable – i.e., their effect should be permanent. Storage devices must play their part in this as well and be reliable and durable given the importance of enterprise data. They must also protect the data, ensure the data is available when needed, and provide fast access to the data. And of course everyone also wants their storage to be inexpensive, adaptable to changing business needs, and easy to use and maintain.

Given the intimate relationship between the Oracle Database and the storage it resides on it is important to ensure the interface between the two is efficient and robust. This is why Oracle and the storage vendors have invested so much in the interface and interrelationship of database and storage. There is a stack of storage management software, programs, and practices available for the Oracle database. All of the components in the stack are highly tuned and optimized for use with the Oracle database. They are the most performant, scalable, reliable, easy to use, and inexpensive way of using storage with the Oracle database.

The Oracle storage solutions stack is organized as depicted below.



The **red** components of the stack represent technology provided in the Oracle database manage storage. **Blue** represents joint Oracle and storage vendor programs to provide storage solutions for Oracle database environments. And **green** represents solution specific best practices like those to achieve high availability. Each layer in the stack addresses things that can compromise data durability. From the database up the storage solutions stack includes the following components.

- **Management and Protection of Oracle Blocks**

- **HARD:** The Hardware Assisted Resilient Data program is a joint offering between Oracle and its storage partners. HARD enables end-to-end validation of Oracle writes and prevents corruptions from being introduced into the I/O path between the database and storage. Not only does the database need reliable storage devices but bugs in the software between the database and the storage should not be allowed to corrupt the database.
- **ASM:** Automatic Storage Management is a cluster-aware volume manager and file system designed for use with the Oracle database to provide the highest levels of performance and reliability while at the same time eliminating manual configuration and I/O tuning sometimes associated with the Oracle database. Being fault tolerant, with automatic I/O tuning and reducing cost – by avoiding the purchase of third party products and the cost of testing and integrating disparate products – make ASM the ideal repository for your Oracle database.

- **OCFS:** Oracle Cluster File System can be used in Linux and Windows environments to host the Oracle home and binaries. OCFS simplifies the management of a Real Applications Cluster (RAC) environment by supporting a single shared Oracle home for all the instances in the RAC cluster.
- **Data Replication for Data Redundancy and Protection**
  - **Data Guard:** Data Guard automates the creation and maintenance of one or more transactionally consistent copies of a production database. The copy, or standby database, is kept synchronized with the production database by the system automatically transferring and applying transactional changes from the production to standby database. It is Oracle's premiere solution for disaster protection and can ensure that the catastrophic loss of the production system will not stop an enterprise from continuing to operate.
  - **Streams:** Streams provides transactional replication for distributed applications. It supports an active-active architecture for multiple distributed databases, heterogeneous platforms as the source and destination, and provides the ability to transform the transactions as they are propagated to the destination database.
  - **OSCP:** The Oracle Storage Compatibility Program was created to validate third party storage solutions like remote mirroring for use with Oracle. Storage vendors provide the ability to capture and propagate changes from one storage device to another. Oracle in partnership with the storage vendors has validated that these technologies work smoothly with the Oracle database and will not corrupt the replica database they maintain.
- **Recovery from Human Errors**
  - **Flashback:** Flashback is an extremely easy to use technology to very quickly undo changes made to the database. Flashback provides fine-grained surgical analysis and repair for localized damage – like when the wrong customer order is deleted. Flashback also allows for correction of more widespread damage yet does it quickly to avoid long downtime – like when all of this month's customer orders have been deleted. Flashback is unique to the Oracle Database 10g and supports recovery at any level of granularity including the row, transaction, table, and database wide.
  - **OSCP:** The Oracle Storage Compatibility Program also validates third party storage solutions like storage snapshots for use with Oracle. Storage vendors and Oracle provide the ability to mirror a volume then split the mirror consistently while the database is running. This provides the ability to take the database volumes back in time to the last snapshot split. Oracle in partnership with the storage vendors has validated that these technologies work smoothly with the Oracle database and will not result in corrupted database files.
- **Recovery from Media Failures**
  - **RMAN:** Oracle Recovery Manager is integrated deep into the database server providing advanced features to efficiently backup and recover the Oracle database. RMAN optimizes performance and space consumption during

backup and integrates with leading tape and storage media management vendor (MMV) products via an RMAN provided API to utilize tape devices and integrate into customers enterprise-wide environment.

- **Oracle Backup and third party MMVs:** Oracle's Backup Solutions Program (BSP) works with vendors to help them produce media management software that integrates with RMAN. To access tape drives and tape libraries, RMAN interfaces with this media management software. Media management software controls tape devices during backup and recovery and manages the loading, labeling and unloading of media, among other functions. An RMAN API enables vendors like Veritas, Legato, or IBM to integrate their media management software with RMAN backup and recovery. New in Oracle Database 10g release 2 Oracle also provides media management software called Oracle Backup. Oracle Backup will be a low-cost means of protecting your database and will reduce complexity by minimizing the need to integrate third party products with the Oracle database. Oracle Backup backs up both the database and non-database files and can you used to protect the entire Oracle environment.
- **Resilient Low-Cost Storage Initiative:** The advent of low cost ATA disk based storage arrays and low cost storage networks together with the introduction of Oracle Database 10g has made it possible to create a database storage grid that has very low cost, excellent performance, and unbreakable availability. The traditional Oracle database configuration consists of a monolithic server connected to a monolithic storage array. Oracle customers have steadily replaced their monolithic SMP servers with RAC clusters of low cost servers running a low cost operating system, Linux. Typically, the costliest component after the host servers is the storage. In many cases, a costly monolithic storage server can be replaced with a grid of low cost storage arrays in a very similar fashion to what is being done with server grids. Oracle in partnership with leading storage vendors has established the Oracle Resilient Low-Cost Storage Initiative to help customers successfully adopt this approach to storage and realize the cost savings grid storage can provide.
- **Best Practices**
  - **ILM:** Information Lifecycle Management (ILM) is concerned with everything that happens to data during its lifetime. New government regulations and guidelines, such as Sarbanes-Oxley, HIPAA, DOD5015.2-STD in the US and the European Data Privacy Directive in the European Union, are a key driving force in how and why data is being retained. Oracle Database 10g is ready today for business ILM. It is capable of storing many different types of data, and storing all of your data in an Oracle database, means that it is much easier to manage, because the data is all in one place, instead of being stored using many different formats. Oracle is now providing a methodology along with the required technology to implement this solution.
  - **MAA:** Oracle Maximum Availability Architecture (MAA) is Oracle's best-practices blueprint based on proven Oracle high availability technologies and recommendations, with the goal of removing the complexity in designing the optimal high availability architecture, and maximizing systems availability. MAA reduces the implementation costs for a highly

available Oracle system by providing detailed configuration guidelines. MAA provides best practices and recovery steps to eliminate or minimize downtime and data loss.

As you can see the Oracle storage stack and components bundled into the Oracle database enable the most effective and efficient relationship between the database and its underlying storage infrastructure. These components are optimized for Oracle and provide the highest possible value to the customer.

**ORACLE**

White Paper Title  
December 2004  
Author: Ron Weiss  
Contributing Authors:

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

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200  
[www.oracle.com](http://www.oracle.com)

Copyright © 2003, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

