

Oracle RAC und das 11.2.0.2 Patch Set – mehr als nur ein Patch Set!?

Markus Michalewicz
Oracle Corporation
Oracle HQ, Redwood Shores, CA, USA

Key words: Oracle RAC, Quality of Service Management, Oracle RAC One Node, ACFS, Oracle Cluster File System, cloud edition, Patch Set 1, Oracle 11.2.0.2

Introduction

Patch Set 1 for Oracle Database 11g Release 2, also called (Oracle Database) 11.2.0.2, is officially a Patch Set. This means, accessibility to the patch is channeled through Oracle Support; no 11.2.0.2 software is available outside of this distribution channel. But even at first glance, 11.2.0.2 is more than just a patch set. Already the overall size of the new patch set indicates that this is not the usual collection of bug fixes. And indeed, the 11.2.0.2 Patch is a complete installation, not just an upgrade, which is a novelty for Oracle, since former patch sets were applied on top of the base version (here 11.2.0.1) typically. Oracle 11.2.0.2 can be used to upgrade or as an independent installation for which no base version is required. This explains the size and lays the ground to introduce new features. It also means that the new patch has to be applied (installed for this matter) in an out-of-place fashion, meaning that a new Oracle Home path needs to be used. This paper, however, will not elaborate on the mechanics of applying 11.2.0.2. It will rather look at the new features introduced with 11.2.0.2, which are often related to Oracle Real Application Clusters (RAC).

Oracle RAC One Node

Oracle RAC One Node is not necessarily a new feature of Oracle Database 11.2.0.2. Oracle RAC One Node was introduced as part of Oracle Database 11g Release 2, but was improved significantly with Oracle Database 11.2.0.2.

Oracle RAC One Node is a new option to Oracle Database 11g Release 2 Enterprise Edition. It provides enhanced high availability for single instance databases, protecting from both planned and unplanned downtime. Oracle RAC One Node provides:

- Always on single-instance database services
- Better consolidation for database servers
- Enhanced server virtualization
- Lower cost development and test platform for full RAC

At least according to the Oracle RAC One Node white paper.

While Oracle RAC One Node was provided via an one off patch with Oracle Database 11g Release 1, enabling the management of an Oracle RAC database as an Oracle RAC One Node database using command line based tools and utilities, it was not properly integrated in the overall management architecture used for Oracle RAC. In addition, Oracle RAC One Node was only available on Linux for Oracle Database 11g Release 2.

These gaps have been closed with Oracle RAC One Node 11.2.0.2. Oracle Universal Installer as well as the DBCA (Database Configuration Assistant) are now fully aware of the concept of an Oracle RAC One Node database and provide options to install just this “one node base” Oracle RAC. This is an improvement over the formerly used patch, which would have converted a full Oracle RAC database into an Oracle RAC One Node database after its install.

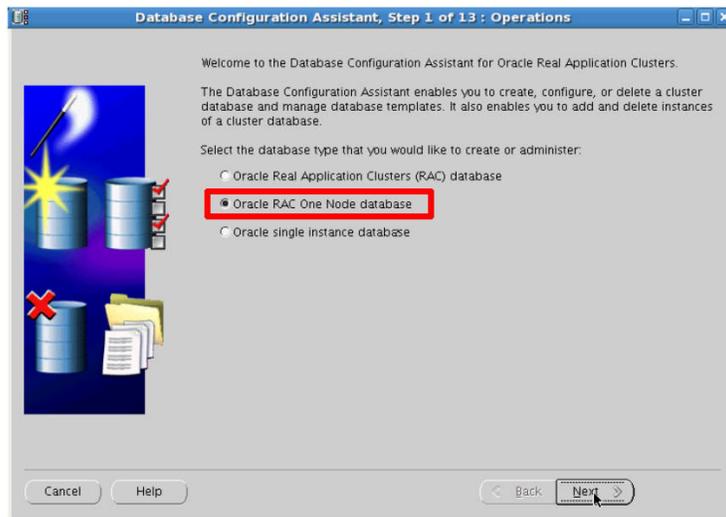


Figure 1 DBCA Oracle RAC One Node Support

Having integrated Oracle RAC One Node into the overall management architecture of Oracle RAC means that Oracle RAC One Node can now be managed using the SRVCTL command line interface. Alternatively, Oracle Enterprise Manager can be used and even advanced features provided by Oracle RAC One Node like Database Relocation (formerly referred to as Omotion) are now fully supported on the command line as well as in the GUI.

Oracle Cluster File System, Cloud Edition (was: Oracle ACFS)

Another feature that was initially introduced with Oracle Database 11g Release 2 (11.2.0.1) was the Oracle ASM based Cluster File System (ACFS). With Oracle Database 11g Release 2, Patch Set 1 (11.2.0.2), ACFS has been renamed to Oracle Cluster File System, Cloud Edition or “Oracle CloudFS”. However, the new name is not the only change.

The “new” Oracle CloudFS comes with a set of new features, while it is now available on all major platforms support for Oracle RAC (Linux, Windows, Solaris, AIX), but not on HP-UX.

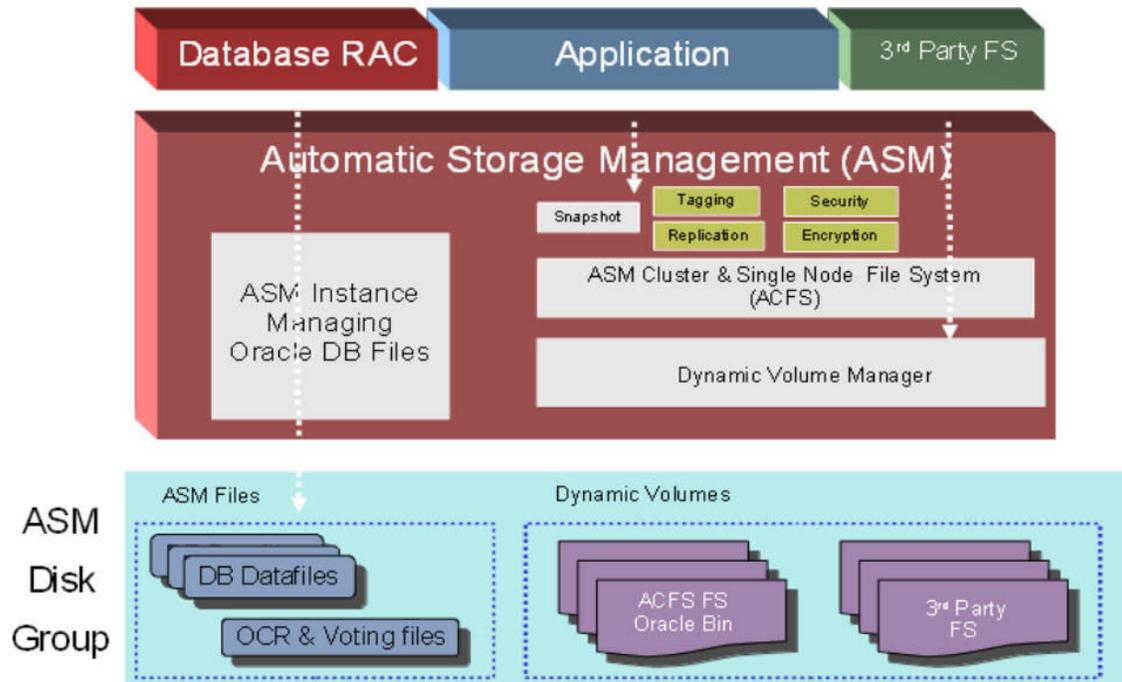


Figure 2 Oracle ASM / ACFS architectural overview

Oracle ACFS also comes with four new features:

- ACFS Tagging
- ACFS Replication
- ACFS Security
- ACFS Encryption

ACFS Tagging

ACFS Tagging is a feature that allows a user to associate one or more files together as a group by assigning a unique 'tag name' attribute. Group operations can be performed based on tagged files that may span across different directories within ACFS file systems within a single node or a cluster. Different groups of tagged files (e.g. 'medical imaging', 'photo album', etc) may be replicated as groups of related files complementing the replication of database files using Oracle Data Guard as a complete DR solution.

ACFS Replication

Oracle ACFS replication is a DR solution that enables replication of Oracle ACFS file systems across the network from a primary to a standby site, providing disaster recovery capability for the file system. Both the primary and standby sites can be standalone or cluster systems. A primary site for one file system can be the standby site for a different file system and vice-versa.

Oracle ACFS replication captures file system changes on the primary site and records the changes in files called replication logs. These logs are transported to the standby site asynchronously where background processes read the logs and apply the changes recorded in the logs to the standby file system. After the changes recorded in a replication log have been successfully applied to the standby file system, the replication log is deleted on both the primary and standby sites.

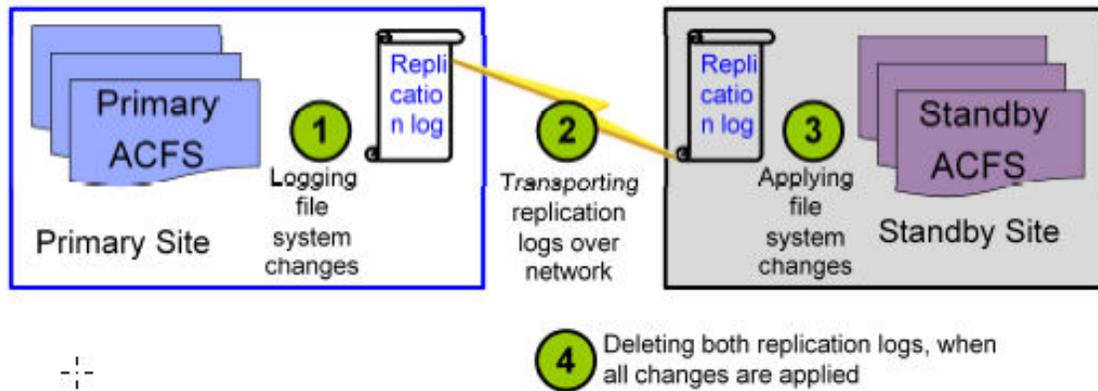


Figure 3 ACFS Replication Overview

ACFS Replication complements Oracle Data Guard and provides the customer with an end-to-end DR solution for all files. You no longer need to rely on 3rd party replication software for external files and Oracle binaries. ACFS Replication is very easy to configure and implement empowering the DBA or the Sys Admin to manage the solution.

ACFS Security

Oracle ACFS Security can be used to protect sensitive data from internal and external threats by providing fine grained access control on top of the access control provided by the operating system.

Oracle ACFS security provides realm-based security for Oracle ACFS file systems, enabling users to create realms and specify security policies for users and groups and therefore control access on file system objects. Oracle ACFS security uses realms, rules, rule sets, and command rules to enforce security policies. The ACFS realm is virtual container of files and directories; access defined by security filters (command rules and rule sets).

An ACFS security administrator manages security for all ACFS file systems in a cluster and has an ACFS Security administration password that is different from his/her OS password. An ACFS Security administration password is stored in an Oracle Wallet for additional security, and all security commands are protected by the realm management password.

Oracle Cluster File System, Cloud Edition

ACFS Encryption

The encryption is another feature that meets regulatory compliance requirements by keeping data in an Oracle ACFS file system in encrypted format and hence protecting against unauthorized use of such data in the case of data loss or theft. Oracle ACFS encryption enables users to encrypt data stored on disk (data-at-rest) and provides secured encryption keys to decode data.

ACFS encryption can be applied to the entire file system or just individual files and directories. It is completely transparent to authorized users, and applications work unchanged with encrypted files. Encrypted and unencrypted files can co-exist on the same file system. System administrators and Oracle ACFS security administrators can initiate encryption operations.

Two keys, the File Encryption Key (FEK) and the Volume Encryption Key (VEK), protect each file. File data is encrypted using a unique FEK; the FEK is stored on disk, and is encrypted using a VEK. VEKs are stored in an Oracle Wallet and can be password protected.

Oracle Database Quality of Service Management

A completely new feature introduced with Oracle Database 11g Release 2 is Oracle's Database Quality of Service Management Feature (QoS). While this feature is officially packaged with the Oracle 11.2.0.2 patch set, it is an Exadata Database Machine exclusive feature, meaning that it can only be used within an Oracle Database Machine as per licensing.



Figure 4 Quality of Service Management introduction

Oracle Database Quality of Service Management allows system administrators to directly manage application service levels hosted on Oracle Exadata Database Machines. Using a policy-based architecture, QoS Management correlates accurate run-time performance and resource metrics, analyzes this data with its expert system to identify bottlenecks, and produces recommended resource adjustments to meet and maintain performance objectives under dynamic load conditions. Should sufficient resources not be available QoS will preserve the more business critical objectives at the expense of the less critical ones.

In conjunction with Cluster Health Monitor, QoS Management's Memory Guard detects nodes that are at risk of failure due to memory over-commitment. It responds by automatically preventing new connections thus preserving existing workloads and restores connectivity once the sufficient memory is again available.

Oracle Database QoS Management is an automated, policy-based product that monitors the workload requests for an entire system. It manages the resources that are shared across applications and adjusts the system configuration to keep the applications running at the performance levels needed by your business. It responds gracefully to changes in system configuration and demand, thus avoiding additional oscillations in the performance levels of your applications.

Oracle Database QoS Management monitors the performance of each work request on a target system. It starts to track a work request from the time a work request requests a connection to the database using a database service. The amount of time required to complete a work request, or the response time (also known as the end-to-end response time, or round-trip time), is the time from when the request for data was initiated and when the data request is completed. By accurately measuring the two components of response time—the time spent using resources and the time spent waiting to use resources—Oracle Database QoS Management can quickly detect bottlenecks in the system; Oracle Database QoS Management then makes suggestions to reallocate resources to relieve a bottleneck, thus preserving or restoring service levels.

Redundant Interconnect Usage

While in previous releases bonding, trunking, teaming, or similar technology was required to make use of redundant network connections between the nodes to be used as redundant, dedicated, private communication channels or “interconnect”, Oracle Clusterware now provides an integrated solution to ensure “Redundant Interconnect Usage”. This functionality is available starting with Oracle Database 11g Release 2, Patch Set One (11.2.0.2).

The Redundant Interconnect Usage feature does not operate on the network interfaces directly. Instead, it is based on a multiple-listening-endpoint architecture, in which a highly available virtual IP (the HAIP) is assigned to each private network (up to a total number of 4 interfaces).

By default, Oracle Real Application Clusters (RAC) software uses all of the HAIP addresses for private network communication, providing load balancing across the set of interfaces identified as the private network. If a private interconnect interface fails or becomes non-communicative, then Oracle Clusterware transparently moves the corresponding HAIP address to one of the remaining functional interfaces.

Oracle RAC Databases, Oracle Automatic Storage Management (clustered ASM), and Oracle Clusterware components such as CSS, OCR, CRS, CTSS, and EVM components employ Redundant Interconnect Usage starting with Oracle Database 11g Release 2, Patch Set One (11.2.0.2). Non-Oracle software and Oracle software not listed above, however, will not be able to benefit from this feature.

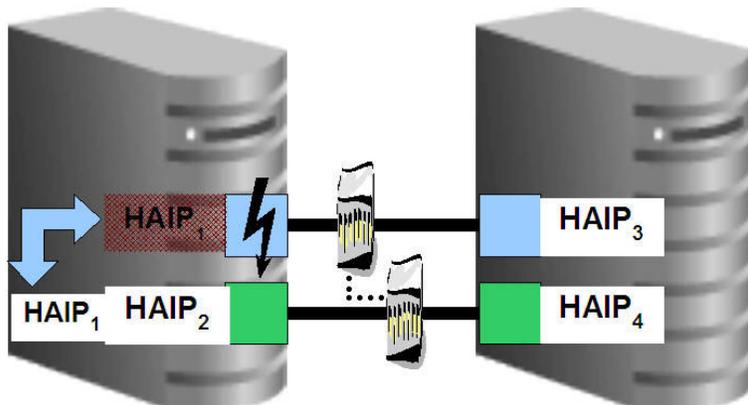


Figure 5 Oracle Redundant Interconnect Usage illustrated

Conclusion

Oracle Database 11.2.0.2 is more than a patch set. Not only can it be used as a full installation, it also introduces some new features, not all of them being optional, as the Redundant Interconnect Usage shows. While those enhancements will eventually help to stabilize the system and improve the service quality provided, they might be changes, which are worthwhile looking at. This paper just introduced the new features.

Markus Michalewicz

Oracle Corporation
500 Oracle Parkway, MS4OP840
USA – Redwood Shores, CA 94065

Telefon: +1(650)5065444
E-Mail: Markus.Michalewicz@oracle.com
Internet: <http://www.oracle.com/goto/rac>