

Taking the 11i/RAC Plunge

Kevin Dahl

Solution Beacon Consulting

This presentation is based on a recent upgrade of a Red Hat 4 Linux based non-RAC 11.5.9 *9i* environment to a Red Hat 4 Linux based 11.5.10.2/10g R2 environment. The focus of this presentation will be an overview of that process with special emphasis on the business justification of RAC, ASM, a split tier configuration for 11.5.10.2 and OATM.

Executive Summary

Migration of a company's existing non-RAC 11.5.9/*9i* environment to 11.5.10.2/10gR2 environment can be a daunting task for a DBA that has not performed a similar migration. This paper provides guidance for that migration based on Oracle Published metalink notes, Oracle documentation, and the experience of performing that migration.

The cost/benefit of RAC will be briefly examined. RAC is primarily a trade-off between increased complexities for the DBA plus additional database license costs versus reduced cost for hardware plus high-availability for the database plus a potential increase in performance. Next, Oracle's relatively new Automatic Storage Management will be explained. This component simplifies disk administration for Oracle RAC databases and can be managed by the DBA. Finally, most of the paper will overview the actual migration process from 11.5.9/*9i* to 11.5.10.2/10gR2. That migration starts with a rapid migration resulting in a temporary single tier implementation. Next, a cluster ready 10gR2 environment is created on 2 64-bit Linux database server including implementation of separate ASM instances. Then the existing database is upgraded from *9i* to 10gR2 and RAC enabled. Finally, parallel concurrent processing is implemented on the middle tier boxes along with a shared application file system.

Overview 11i 10g R2 RAC from 11i 9i

10g R2 RAC cost/benefit analysis

The typical costs observed in Oracle RAC customers are:

- Increased time and complexity of implementation and subsequent maintenance
 - This included CRS implementation/administration. The bulk of this "cost" is the initial install of CRS
 - ASM implementation/administration. The major portion of this additional "cost" is the initial configuration and setup of ASM including the creation of a separate ASM instance for administration of ASM
 - Other DBA tasks – These would be the extra steps a DBA would go through to shutdown and startup and environment, additional effort during cloning and additional effort during patching
- Additional machines to administer
 - At a minimum, the dba tier hardware would double. That means twice the number of boxes to maintain from an OS perspective as well as a hardware perspective
- Increased license cost
 - This potential cost can be verified with your Oracle Sales representative

The typical benefits observed in Oracle RAC customers are:

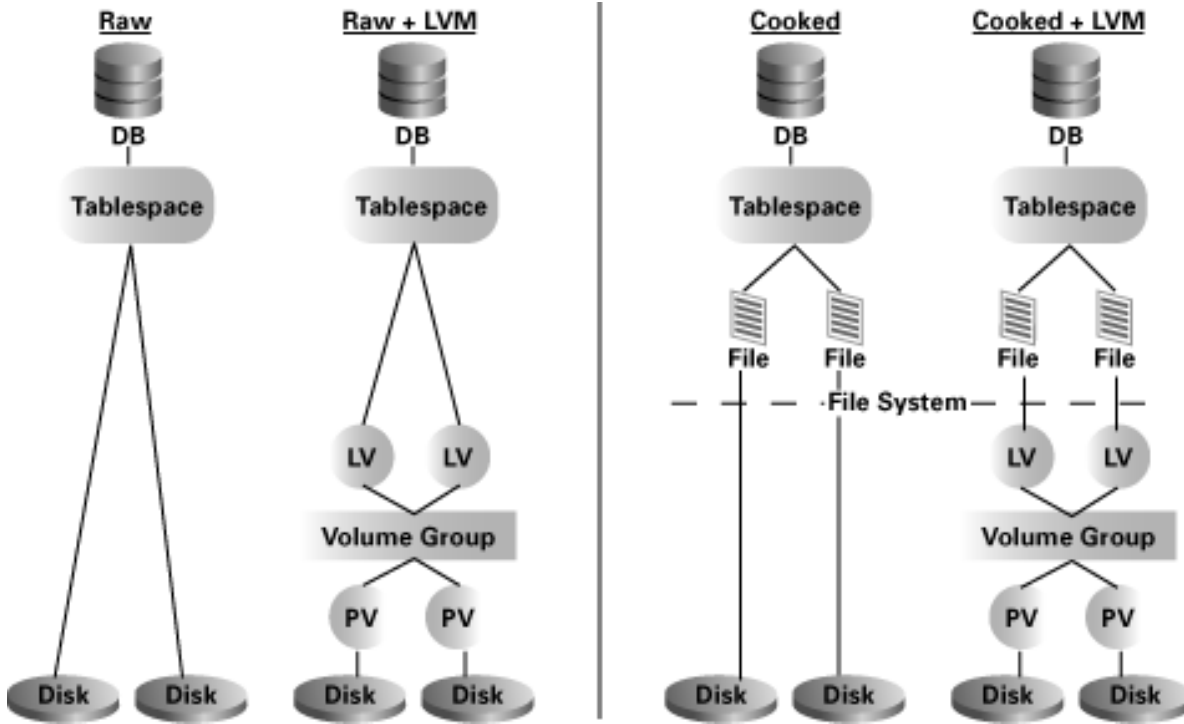
- Non Disaster/Recovery high availability – RAC will provide non TAF for oracle applications
- Increased performance – the additional node(s) provide additional processing power
- Depending on user base, would be able to use "commodity" hardware and operating system. Currently "commodity" RAC nodes would be a 4 CPU box with 32 gig memory. Generally a couple of those boxes and a couple of middle tier boxes would provide a sufficient hardware environment for around 200 concurrent users

Overview of ASM

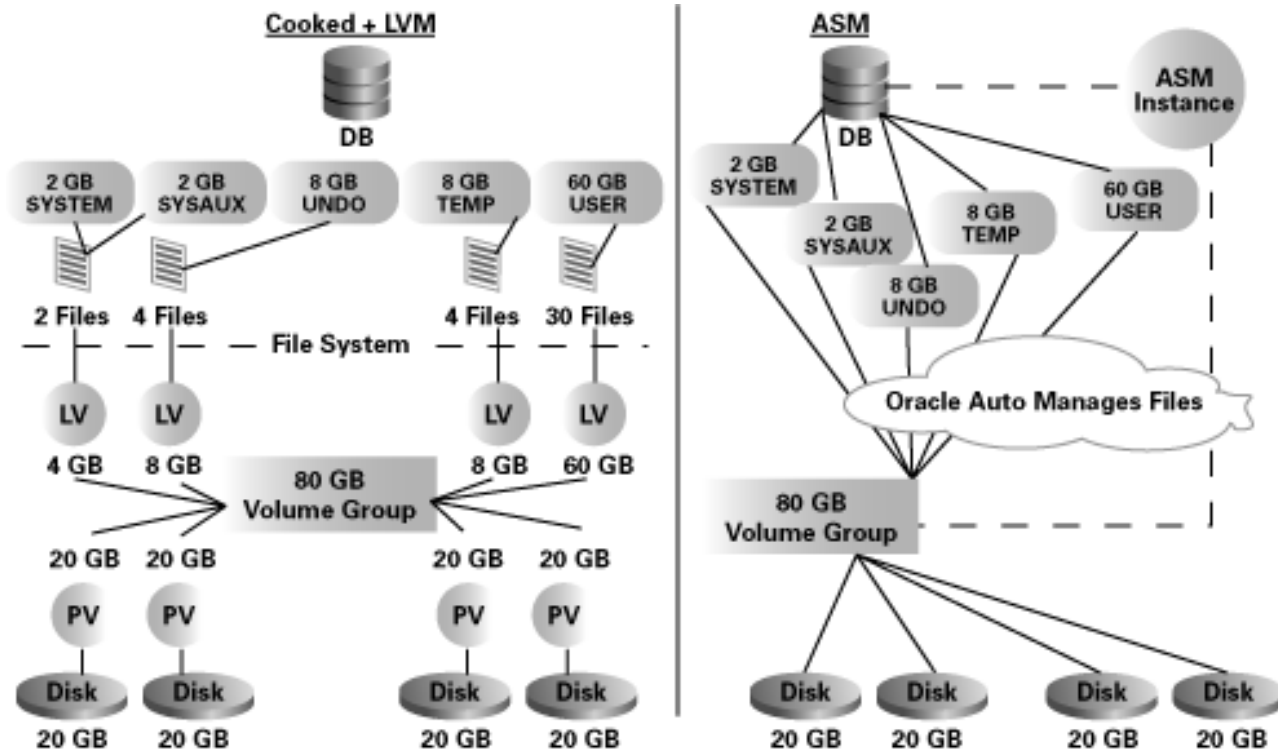
The Oracle Automatic Storage Management (ASM) feature of Oracle Database 10g provides a virtualization layer between the database and storage so that multiple disks can be treated as a single disk group and disks can be dynamically added or removed while keeping databases online. Existing data will automatically be spread across available disks for performance and utilization optimization. As a vertically integrated file system and volume manager, purpose-built for Oracle database files, ASM provides the

performance of async I/O with the easy management of a file system. ASM provides capability that saves the DBAs time and provides flexibility to manage a dynamic database environment with increased efficiency.

Overview of ASM-traditional non ASM storage



Overview of ASM-cooked vs. ASM



To summarize, ASM provides the following benefits:

- Simple storage management interface
- The disk storage is “virtualized” into a disk group. Virtualized means that underlying storage may consist of multiple disks or portion of those physical disks but are presented as a single ASM disk
- ASM automates the placement of database files within those disk groups. The following tablespace creation statement will illustrate that (disk01 is the ASM diskgroup):

```
CREATE TABLESPACE APPS_TS_TX_DATA DATAFILE '+disk01' SIZE 8000M EXTENT MANAGEMENT LOCAL
AUTOALLOCATE SEGMENT SPACE MANAGEMENT AUTO;
```

- ASM spreads data evenly across all available storage resources to optimize performance and utilization. This even distribution of database files makes the manual I/O performance tuning obsolete
- ASM provides three mirroring options for protection against disk failure: none, two-way, and three-way mirroring. The none option allows ASM to defer protection to the storage device if so desired
- ASM enables the DBA to change the storage configuration without having to take the database offline
- ASM automatically rebalances files across the disk group after disks have been added or dropped

High-level process of upgrading to 11i/10gR2 RAC from 11.5.9/9i

The remaining sections present an overview of steps used to upgrade an existing 11.5.9/9i to 11.5.10.2/10gR2 RAC

Upgrade existing 11.5.9/9i to 11.5.10.2/9i via SB rapid migration technique

Solution Beacon’s rapid migration approach was used to upgrade the existing 11.5.9 applications environment to 11.5.10.2. This was done as a single tier install on one of the two middle tier boxes. These are the high-level steps for that process:

- Perform fresh install of 11.5.10.2 software

- Upgrade the 11.5.10.2 oracle home from 9.2.0.6 to 9.2.0.7
- Replace the installed database with the cloned production database
- Applied pre 3480000 patches
- Applied patch 3480000 – this is the 11.5.10.2 maintenance patch
- Applied post 3480000

Later in the process, this will become a split tier install. That means the database ONLY will be moved to the RAC nodes and upgraded to 10gR2. The oracle home on the middle tier will be deleted. All other application processing will remain on the middle tier nodes.

Install Oracle Clusterware 10.2.0.1

Oracle clusterware is a portable cluster management solution that is integrated with the Oracle database. Note that a cluster is made up of multiple interconnected computers or servers that appear as if they are one server to end users and applications. Oracle Clusterware is also a required component for using RAC. RAC uses Oracle Clusterware for the infrastructure to bind multiple servers so that they operate as a single system.

These are Oracle Clusterware requirements:

- Voting disk – records node membership. This is created on a shared raw device
- Oracle Cluster Registry (OCR) - records cluster configuration information. is created on a shared raw device
- Private interconnect – private network used by the nodes to communicate with each other. It should use Datagram Protocol (UDP) on a Gigabit Ethernet. This interconnect is the “glue” of the RAC nodes

The following are the major clusterware services

Cluster Synchronization Services (CSS) —Manages the cluster configuration by controlling which nodes are members of the cluster and by notifying members when a node joins or leaves the cluster

Cluster Ready Services (CRS) —Program for managing high availability operations within a cluster. A cluster resource which could be a database, an instance, a service, a Listener, a virtual IP (VIP) address, an application process, and so on

Event Management (EVM) —A background process that publishes events that crs creates

Oracle Notification Service (ONS) —A publish and subscribe service for communicating Fast Application Notification (FAN) events

RACG —Extends clusterware to support Oracle-specific requirements and complex resources. Runs server callout scripts when Fast Application Notification (FAN) events occur. (Applications use FAN to enable rapid failure detection, balancing of connection pools after failures, and re-balancing of connection pools when failed components are repaired. The FAN notification process uses system events that Oracle publishes when cluster servers become unreachable or if network interfaces fail)

Process Monitor Daemon (OPROCD)—This process is locked in memory to monitor the cluster and provide I/O fencing

Install Oracle Database Software 10.2.0.1

This is the first step in the installation of the 64-bit 10gR2 software on the linux boxes reserved for the database software. This oracle home will replace the one that was installed on the middle tier box at the beginning of this upgrade. This RDBMS software installation is a typical database software install using the gui runinstaller.

Install Database Components from Oracle 10g (10.2.0.1) Companion CD

The next step in the installation of the 10.2.0.1 software is to include components from the companion CD. When prompted, choose Oracle Database 10g Products 10.2.0.1. Those products include the following:

- Oracle JDBC Development Drivers
- Oracle SQLJ
- Oracle Database Examples
- Oracle Text Supplied Knowledge Bases
- Oracle Workflow Server
- Oracle Ultra Search
- Natively Compiled Java Libraries
- JPublisher

Upgrade Cluster Ready services , database software to 10.2.0.2

The final step in the installation of 10gR2 software is to upgrade the base 10.2.0.1 version to version 10.2.0.2. Again, this is done via the gui runinstaller. Note that you must shutdown cluster ready services first before launching the install.

Upgrade Oracle9i (9.2.0.6) Database to 10g (10.2.0.2)

Recall the the database currently resides on one of the middle tier boxes. The database files are now NFS mounted on the primary RAC node retaining the same file path as on the source. Next, the manual upgrade steps described in the 10g upgrade manual are followed. Here is a summary of those steps:

- Run @utlu102i.sql provided in 10.2.0.2 \$ORACLE_HOME/ rdbms /admin against 9i. Analyses and provides list of tasks
- Adjust init.ora for 10g
- With in sqlplus as sysdba , STARTUP UPGRADE
- Run catupgrade.sql to upgrade the database
- Run utlu102s.sql to see what has been upgraded
- Compile invalid objects

Create ASM Instances/Diskgroups

The first step in creating ASM diskgroups is to install/configure ASM support library, ASMLib . ASMLib is an application program interface (API) developed by Oracle to simplify the operating system-to-database interface, and to exploit the capabilities and strengths of vendors' storage arrays. The purpose of ASMLib, which is an optional add-on to ASM, is to provide an alternative interface for the ASM-enabled kernel to discover and access block devices. To download ASMLib for Linux, go to <http://www.oracle.com/technology/tech/linux/asmlib>

After installing ASMLib, use fdisk to make raw storage available to ASM. Then, create an ASM disk via oracleasm. That ASM diskgroup will be detected by the installer when the ASM instance is next created. Although the next task is not required, it is highly recommended to create a separate ASM instance instead of relying on ASM components within an existing database. A separate ASM instance allows independence from a database and can be used for multiple databases on a given box. To install ASM, first install a separate ASM oracle home (software only). Then, create an ASM instance via dbca. Ensure the following selections are made:

- On the welcome screen, select Oracle Real Application Clusters database
- On the Step 1 of 4:Operations screen, Select Configure Automatic Storage Management
- On the Step 2 of 4:Operations screen, Select all cluster nodes that appear
- On the Step 3 of 4:Operations screen, enter the system password and use an ifile rather than spfile
- On the Step 3 of 3:ASM Disk Groups screen, click the mount command button so the detected disk group is mounted

The result is a new ASM instance that will provide a managed ASM diskgroup for the migrated database.

Convert to OATM

This point in the process was a convenient place to do the OATM conversion. The target ASM datafiles created in step 2 of the menu was a convenient combination of a conversion to both OATM and ASM. \$FND_TOP/bin/fndtsmig.pl is the OATM migration script and it presents the following menu, which outlines the OATM process.

1. Migration Sizing Reports
2. Create New Tablespace
3. Generate Migration Commands
4. Execute Migration Commands
5. Run Migration Status Reports
6. Run Post Migration Steps
7. Run Customization Steps
8. Run Migration in Batch Mode

Note that a sub-menu selection under option 1 indicates how large the new tablespaces need to be. Also note the migration status report under selection 5 will need to run until the migration is 100% complete. The OATM migration will only handle Oracle

Application tablespaces. Remaining tablespaces will have to be examined to determine if they should be dropped or moved to ASM storage under their same names.

Convert 10g Database to RAC

RCONFIG was not used to convert the database to RAC. Rather, the conversion was done manually. The following were key steps in that conversion:

- Ensure a new appsutil.zip file is generated on the applications tier and transferred to the primary 10g db tier \$ORACLE_HOME. Then, execute autoconfig on the 10g db tier
- On the first RAC node
 - Execute adpreclone.pl on the primary 10g db tier. This will gather db info on the existing node and create a clone subdirectory
 - Ensure the database and database listener are shutdown
 - Ensure the ASM instance(s) are still running
 - Execute adcfgclone.pl. Ensure you respond that this is the first RAC node, there are 2 instances, the service name and the port pool, correct hostnames and instance number. Also, there will only be 1 datatop because ASM is being used
 - Zip the appsutil directory in the 10g \$ORACLE_HOME
- On the second RAC node
 - Copy the zipped appsutil directory from the first RAC node and unzip it in the \$ORACLE_HOME on the second node
 - Run adcfgclone.pl on the second node. Ensure you respond that this is RAC, this Not the first RAC node, the appropriate hostname, sid
- Edit listener.ora and tnsnames.ora to accommodate RAC
- Create second undo for the second instance
- ensure cluster tags are correct in xml files. For example, s_dbCluster must be set to true
- run autoconfig on second tier
- shutdown db
- Create spfile from pfile ;
- startup database instance

Establish Applications Environment for RAC

The following steps were used to establish the applications environment for RAC

1. The Applications environment was sourced
2. The tnsnames.ora from node 1 was copied to the 806 tnsnames.ora
3. Test connectivity to one of the instances in the RAC environment before running AutoConfig in the next step
4. Execute AutoConfig by running the command:
\$AD_TOP/bin/adconfig.sh contextfile=\$APPL_TOP/admin/<context_file>
5. Check the \$APPL_TOP/admin/<context_name>/log/<MMDDhhmm> AutoConfig log file for errors
6. Source the environment by using the latest environment file generated
7. Verify the tnsnames.ora and listener.ora files located in the 8.0.6 ORACLE_HOME, and \$IAS_ORACLE_HOME/network/admin. Ensure that the correct TNS aliases have been generated for load balance and fail over. Ensure that all the aliases are defined using the virtual hostnames
8. Verify the dbc file located at \$FND_SECURE. Ensure that the parameter APPS_JDBC_URL is configured with all instances in the environment and load_balance is set to ON

Configure Parallel Concurrent Processing

Parallel concurrent processing takes advantage of multiple middle tier boxes. The processing of concurrent request is spread across multiple machines when parallel concurrent processing is configured. The following steps were used to setup parallel concurrent processing:

1. setup PCP
 - a. FNDSM configured and working on all nodes
 - b. Internal monitor configured and working all nodes
2. Setup transaction managers
 - a. Must edit RDBMS ifile pointed to by init.ora adding _lm_global_posts =TRUE and _immediate_commit_propagation =TRUE

- b. In apps, change the profile option 'Concurrent: TM Transport Type' to 'QUEUE' and setup secondary and primary nodes names for TMs
 3. Load balance concurrent processing tiers
 - a. Edit context xml file with CP load balance alias and run autoconfig

High-level process of upgrading to 11i/10g RAC: Conversion to shared application file system

To enable shared application file system:

On the existing middle tier node with the application software

1. \$ cd \$COMMON_TOP/admin/scripts/<CONTEXT NAME>
2. \$ adstpall.sh apps/apps
3. \$ cd \$FND_TOP/patch/115/bin
4. \$ perl -I \$AU_TOP/perl txkSOHM.pl
5. NFS share of the above application file system
6. Run preclone on existing system

On the new middle tier node

7. \$ cd <AD_TOP>/bin
 \$ perl adclonectx.pl sharedappltop contextfile=<Applications Context File for the existing node>
 \$ cd <FND_TOP>/patch/115/bin
 \$ perl -I <AU_TOP>/perl txkSOHM.pl

See metalink note: *Sharing the Application Tier File System in Oracle Applications 11i - 233428.1* for a description of the prompts that will be shown during the execution of the txkSOHM.pl perl script

Conclusion

With any 11i migration multiple migrations must be performed. It is recommended that at least 3 test iterations are performed. The first iteration would be for technical familiarity. The next iteration would be to capture timings and refine the process. The third iteration is to further refine the process and ensure a repeatable process now exists for implementation into production during the proposed go-live window.

The information presented in this paper is taken from the following sources. These are the documents that provided the roadmap during the actual conversion.

- *Configuring Oracle Applications Release 11i with 10g R2 RAC and ASM* – metalink note 362135.1
- *Configuring Oracle Applications Release 11i with 10g RAC and 10g ASM* – metalink note 312731.1
- *Oracle® Applications System Administrator's Guide - Configuration Release 11i* Part No. B13925-06
- *Advanced Configurations and Topologies for Enterprise Deployments of E-Business Suite 11i* metalink note 217368.1
- *Using Oracle Applications with a Split Configuration Database Tier on Oracle 10g Release 2* – metalink note 369693.1
- *Shared APPL_TOP FAQ* – metalink note 243880.1
- *Sharing the Application Tier File System in Oracle Applications 11i* - metalink note 233428.1

About the Author

Kevin Dahl has been working with Oracle Applications since 1991. He worked in insurance, telecommunications, and energy industries before entering consulting. His primary function with Solution Beacon is to provide Oracle Application DBA consulting to customers. That involves new implementations, upgrades, performance tuning and other database related tasks. Mr. Dahl holds a Masters degree in Computer Systems Management from Creighton University.