

# HP ProLiant servers running Oracle 10g Real Application Clusters (RAC) on Windows Server 2003 (32-bit)

[REVISED – 11 JUL 2005]



Overview .....	2
Assumptions .....	2
References .....	2
HP-recommended architectures for Oracle RAC on Windows 32-bit .....	2
Server recommendations: .....	3
Storage recommendations: .....	3
Redundancy options: .....	3
Installation example: .....	4
Pre-requisites for installation of Oracle software .....	4
Windows Server 2003 Enterprise Edition 32-bit installation .....	4
Network configuration (to be performed on each node) .....	4
HP software installation .....	5
External shared storage configuration .....	5
Configuration of External shared storage for Oracle software .....	5
Configuration of HP StorageWorks MSA1000 storage or MSA500 storage .....	7
Configuration of StorageWorks EVA 3000/5000 storage .....	7
Configure shared storage in Windows Server 2003 (see RAC Guide) .....	8
Oracle Cluster Ready Services (CRS) software installation .....	9
Oracle 10g Database Enterprise Edition software installation .....	9
Post-installation configuration .....	9
Apply Patchset 10.1.0.4 Patch Set 2 for 32-Bit Windows Server 2003 .....	9
Creation of Sample Oracle 10g Database .....	10

## Overview

This document describes the HP-recommended configuration and installation details for Oracle® 10g Real Application Clusters (RAC) to leverage 32-bit HP ProLiant servers in a Windows® Server 2003 environment. This document has two objectives: 1) define a set of HP recommended architectures for use in deploying Oracle RAC that balances the high availability and scalability characteristics of Oracle RAC with an HP platform of compatible and complementary characteristics, and 2) provide details on an example 10g RAC installation process that highlights the configuration requirements and options specific to the Windows Server 2003 32-bit environment.

## Assumptions

It is assumed that the reader is familiar with configuring computer networks and installing software in a Microsoft Windows Server 2003 environment and that the Hewlett Packard (HP) and Oracle documentation for Oracle 10g Release 1 Database and Real Application Clusters (RAC) for 32-bit Windows are available as references. Also, it is assumed that the hardware configuration has been completed and the necessary software is available.

## References

The following documents were used during the development of this installation guide, with many of the procedural steps derived from them. A shortcut name for each document is indicated, which is used in the procedures as a reference.

1. Oracle Real Application Clusters, Installation and Configuration Guide, 10g Release 1 (10.1.0.2.0), referred to as "RAC Guide"
2. Oracle Database Release Notes, 10g Release 1 (10.1..0.2.0) for Windows, referred to as "DB Notes"
3. Oracle Database Installation Guide, 10g Release 1 (10.1.0.2.0) for Windows, referred to as "DB Install"
4. Oracle Database Patch Set Notes, 10g Release 1 (10.1.0.4) Patch Set 2 for Microsoft Windows (32-Bit), referred to as "PatchSet"

## HP-recommended architectures for Oracle RAC on Windows 32-bit

Oracle Real Application Clusters (RAC) is a shared-everything architecture in which all nodes in the cluster have access to the data on shared disk. Data does not need to be partitioned among the nodes and all data synchronization and cache coherency are controlled by the Oracle software. This model enables seamless migration of any Oracle database into a RAC environment. Once enabled, the new RAC system immediately provides significant high availability capabilities – such as clients being buffered from any server or instance level failure. From this point, many additional Oracle and third-party tools that provide performance and scalability can now be applied. Since there are now multiple nodes capable of processing any database function, opportunities are opened to prioritize access by process, user, data element, etc. or to expand processing capacity at will by adding new nodes, which again, increases availability. To fully leverage this architecture, a

storage environment that is equally flexible and scalable is a must. Traditional SCSI environments are limited to two-node access and are catered to smaller environments that are not expected to expand or are used strictly as application development platforms. NAS and iSCSI are unproven options that currently bring more challenges than value. HP is working to evolve these options for future use with Oracle RAC. For these reasons, HP recommends that current deployment of RAC for production purposes be on Fibre Channel SANs.

From a server perspective, RAC offers an option to the traditional SMP scale-up model because all nodes in the cluster have concurrent access to a common database. Therefore, expanding processing capacity is accomplished by adding more processors and memory to the existing cluster nodes or by adding additional nodes. A well designed client interface to the RAC cluster that leverages Oracle and third-party technologies will seamlessly integrate additional nodes, rebalancing workload dynamically or statically for short term needs. This model offers significant total-cost-of-ownership advantages including much lower initial costs, closer alignment of capacity needs and costs, and the investment protection from 'fork-lift' upgrades. For these reasons, two- and four-processor server models are the focus of HP engineering work around RAC.

RAC customers are almost always interested in the high availability features, as well as the potential for scale-out growth. For this reason, HP recommends redundancy and recovery options for storage and the RAC cluster interconnect to mitigate these components as potential points of system failure. It is important to note that a potential alternative to redundancy at this level is a larger node count, which effectively provides more paths to storage and cluster nodes without the added complexity of additional software.

In summary, HP recommends the following high level architectural approaches to RAC on 32-bit Windows-based ProLiant platforms:

#### Server recommendations:

- HP ProLiant DL580 G2 or later server with 1 to 4 Intel Xeon processors (any speed)
- HP ProLiant DL380 G2 or later server with 1 to 2 Intel Xeon processors (any speed)
- HP ProLiant BL20p Blade servers with 1 to 2 Intel Xeon processors (any speed)
- HP ProLiant DL380 G3 or later Packaged Cluster with MSA500 storage

#### Storage recommendations:

- HP StorageWorks MSA1000 Fibre Channel SAN with 2GB switch
- HP StorageWorks EVA3000 / 5000 Fibre Channel SAN with 2GB switch
- HP Fibre Channel HBA FCA2204 or FCA2204DC

#### Redundancy options:

- HP Secure Path for Windows - IO multi-path management software
- Microsoft MPIO multi-path environment (only Basic version available now; Full function version available in future from HP)
- NIC Teaming drivers from HP

Oracle Real Application Clusters one-time initial deployment can be complex and time consuming, with the various lengthy Oracle documents describing the requirements and procedures. This white paper provides an example of the installation procedure, based on the Oracle documentation and actual experience with installing 10g RAC on the 32-bit HP server and storage platforms.

## Installation example:

A "step-by-step" procedure is described below for installing a 2-node Oracle 10g Real Applications Cluster (RAC) using OCFS for shared Oracle home binary files and ASM for Oracle database files on HP ProLiant servers running Windows Server 2003 Enterprise Edition (32-Bit). Note that the sample database is created for testing and verification of the system. The procedure is organized into the following four sections:

- Pre-requisites for Installation of Oracle software
- Oracle 10g Database Enterprise Edition software Installation
- Post-installation Configuration
- Creation of Sample Oracle 10g Database.

## Pre-requisites for installation of Oracle software

### Windows Server 2003 Enterprise Edition 32-bit installation

1. Install the HP ProLiant server nodes, network connections, and shared storage EVA 3000/5000 or MSA1000 or MSA500 array hardware, setting up any redundant network (teaming) and or redundant storage paths, as desired.
2. Use HP's Smart Start CD software to boot up each server and configure the local storage array and then follow the Setup OS steps, using the Microsoft Windows Server 2003 Enterprise edition software for 32-bit support. If the server comes with Windows OS already pre-installed, then complete the installation steps as indicated in the documentation that comes with the server.

### Network configuration (to be performed on each node)

1. Each node in the cluster must have at least two network interfaces, one for the public network and one for the private cluster interconnect. These need to be on separate networks. After editing the hosts file on each server node with the proper IP addresses and node name (<node name> for public address; <node name>\_pvt for private address, and <node name>\_vip for the virtual address needed for 10g), verify that each node can be pinged from either the public or private node names (see RAC Guide).

Note: The hosts file is located in the <Windows home directory> \system32\drivers\etc folder.

2. Remember to check the Advanced tab in Network and Dial-up Connections and make sure the Public LAN connection is listed first, followed by the private cluster interconnect (as described in RAC Guide).
3. Oracle recommends (see RAC Guide) that the environmental variables for TEMP and TMP point to hard drive with at least 100MB of free space. The drive location associated with

the temp space can be displayed by right-clicking My Computer, select Properties, select Advanced tab, and click the Environment Variables button. In the User variables section, edit the two entries for TEMP and TMP. If the drive does not have enough free space, then edit the variable to point to a different drive that does.

4. If NIC teaming is desired for the private cluster interconnect or both the private and the public network ports, you must use the HP utility to do this. Usually, this HP utility is already installed from the HP Smart Start CD, but if not, you can install it from this CD by selecting the Software tab to run the ProLiant Support Pack Custom installation and then select the HP Network Configuration utility and install it. Once installed, you may access the utility from the icon in the lower right corner of the desktop, near the time display. Select the pair of NICs to be teamed and choose Fault Tolerant and Fail on Fault. Assign an IP address to the team, using the Windows Network utility and update the hosts file with the new team IP address.

## HP software installation

5. Install other HP software, such as the external shared storage firmware and files for Secure Path or Windows MPIO redundant storage solutions. The software utility and instructions for flashing the MSA1000 firmware (version must be 4.24 or later) or the EVA3000/5000 firmware (version must be 3.010 or later) can be obtained from the HP website. Updating firmware is done from one node in the cluster to the storage controllers.
6. There are two redundant storage solutions available for the ProLiant servers: HP Secure Path's multipath solution and Windows MPIO solution. At this time, both are supported for the EVA 3000/5000 storage and the MSA1000 storage, but the MPIO solution has less functionality (Basic version). Only one of these solutions may be installed on the servers—do not install both. The installation of Secure Path or MPIO files from HP needs to be done on each server node in the cluster. The HBA driver should already be installed during the HP Smart Start installation of Windows process. If Secure Path or MPIO are required, either multi-path software must be installed now. Obtain either the HP Secure Path software or the Microsoft MPIO software from the HP website, choosing the version that is appropriate for Windows Server 2003 32-bit. Another multipath solution is available for the Packaged cluster's MSA500 and this software comes with the cluster hardware, along with the multipath option kit.
7. After installation of Secure Path, MPIO, or MSA500 multipath, you may connect the second storage path (controller, cable, HBA port, and switch) to the MSA1000 or EVA 3000/5000 or the Packaged cluster MSA500.

## External shared storage configuration

### Configuration of External shared storage for Oracle software

1. There are several choices available for configuring the Oracle home binary software and the Oracle database data files in a Windows Server 2003 environment. The choices made determine how to configure the shared storage—both the number and size of logical drives. One choice is to install the Oracle 10g Enterprise database binaries (Oracle Home) either on each of the local storage drives for each server node in the RAC cluster or once on the shared storage drives using Oracle Cluster

File System (OCFS). Another choice is to use either OCFS or Automatic Storage Management (ASM) to support the Oracle database data files on the shared storage. A third option to use raw devices is also available, but the recommendations from Oracle are to use OCFS for the Oracle home and ASM for the data files. These options and the logical device configuration choices for each are documented in the Oracle documentation (see RAC Guide, chapters 7 and 9).

2. The recommended number and size of logical drives for each choice, as of the writing of this paper, are as follows:
  - Oracle Home binary software using OCFS on shared storage requires a minimum of one logical drive of at least 3GB
  - OCFS for Oracle database data files requires a minimum of one logical drive of at least 3GB, and if separate recovery area is used, another drive of 3GB
  - ASM for Oracle database data files requires a minimum of one logical drive of 1.15GB with External redundancy, 2 drives of 2.3GB for Normal redundancy (the default), and 3 drives of 3.45GB for High redundancy. If separate recovery files are to be used, then additional space of two times is needed on the same logical drive—or a separate logical drive can be created for this purpose.
  - The two files for CRS (100MB and 20MB) can be installed on the OCFS data files volume, if that is used; otherwise, two logical drives are required on the shared storage for this purpose.

The table below summarizes these choices:

	OCFS	ASM	Raw devices
Oracle Home binary	3 GB	Not applicable	See Oracle document
Oracle Data files	3 GB	1.15 GB – 3.45 GB, depends on redundancy	See Oracle document
Oracle Recovery area	3 GB	1.15 GB – 3.45 GB, depends on redundancy	See Oracle document
CRS files	Included	Not applicable	Voting – 20 MB OCR – 100 MB

3. Oracle requires that the Windows automount feature is enabled on all server nodes before installing Real Application Clusters (see RAC Guide). To check the current setting, open a command window and from the prompt, execute the following commands:

```
C:\> diskpart
```

```
DISKPART > automount
```

If the message displayed indicates that automatic mounting of new volumes is disabled, then execute the following command:

```
DISKPART > automount enable
```

and the message should indicate that automount is enabled. Type exit to quit. This must be done on each of the nodes in the cluster before installing Oracle.

## Configuration of HP StorageWorks MSA1000 storage or MSA500 storage

1. Confirm that the shared storage (MSA1000) is powered up. Use HP's built-in ACU utility on Node 1 to configure the MSA1000 storage to have the number and size of logical drives required to support the desired configuration, as described previously. Access the ACU tool from Start->Programs->HP System Tools->HP Array Configuration Utility->HP Array Configuration Utility. Select the MSA1000 controller in the Configuration view.
2. If there is a pre-existing array configuration displayed, clear the configuration by performing the following:
  - Click Clear Configuration under the Configuration View. A confirmation window is displayed.
  - Click OK in the confirmation window.
  - Click Save. A confirmation window is displayed.
  - Click OK in the confirmation window.
3. Click Create Array under the Select a task for heading and a list of all drives available is displayed. Configure the storage array by performing the following:
  - Select all the disks and click OK. Select the newly created Array and select Create Logical Drive.
  - Select RAID Level 1+0 and enter the recommended size and take the defaults for other fields.
  - Repeat the steps to create a logical drive for as many drives as are needed.
  - Click Save. A confirmation window displays. Click OK to save the configuration.
4. Click Selective Storage Presentation (SSP) from the right windowpane. Configure SSP by performing the following:
  - Select Enable and OK.
  - Select Windows from the Host Mode drop box. (Scroll right to see these options)
  - Select all the logical drives for each adapter ID that need access to this server. (For this installation process select all of them).
  - Click OK. A warning window is displayed. Select OK in this warning window.
5. Click Save. A confirmation window is displayed. Select OK to save the configuration.
6. Click Exit ACU. The ACU utility closes.
7. Repeat steps 4 through 6 on each of the other nodes in the cluster, to confirm that the shared storage is visible and that SSP is configured.

## Configuration of StorageWorks EVA 3000/5000 storage

1. Obtain hostname, IP address, and HBA WorldWide ID (WWID) for each server in the cluster. To find the HBA WWID, you must use the Emulex lputilnt utility. If it is not already installed, you can obtain it from the HP website, as part of the software available for the EVA 3000/5000 storage system. Once installed, you can start lputilnt by choosing Start->Programs->Emulex. After the lputilnt utility screen is displayed, click the HBA listed in the left pane of the window that is displayed. Choose Configuration Data as the

- category and scroll down to 16-World Wide Name in Region to see the WWID for that HBA. If you have two HBA ports, there will be two HBA's listed in the left pane.
2. Using the HP SAN Appliance and the Command View EVA utility and configure the EVA host profiles for the server nodes in the cluster. You will need to Add each host, and enter the hostname, the IP address (if not using DHCP), and from the pull-down menu, select the WWID for the HBA in that server. Select Windows as the OS. If you have more than one HBA port per server, select Add port and select the additional WWID for that server's other HBA port. You should have a host profile for each of the server nodes in the cluster.
  3. Also using the HP SAN Appliance and the Command View EVA utility configure a LUN for the Oracle database on OCFS or ASM and present it to the server nodes in the cluster. Start by creating a Vdisk—you can take the default Vdisk name and select the desired disk group from the Disk Group Name drop down list. Select the Vraid1 option for redundancy and enter the recommended size in the size text field. Create as many additional Vdisks as are needed for the chosen configuration (OCFS or ASM for data and shared Oracle Home binary files).
  4. Under the Vdisk created in the previous step, click the Active icon. Click the Presentation tab and then click Present. Select all of the server hosts for which profiles were created in step 2 above.
  5. Reboot each of the server nodes in the cluster.

#### Configure shared storage in Windows Server 2003 (see RAC Guide)

1. On one node in the cluster, open the Windows Server 2003 Computer Management utility and select Disk Management under Storage.
  2. The Windows Server 2003 Initialize and Convert disk wizard is displayed to initialize the disk volume(s) and the option to convert from Basic disk to Dynamic disk is offered. Initialize the disk, but do not convert to Dynamic disk on the next screen, as this is not supported for an Oracle database. If it is necessary to change a Dynamic disk back to a Basic disk, all the data on the disk will be lost. Confirm that the disk is Basic and Master Boot Record (MBR) by right-clicking the disk and choosing Properties. Select the Volumes tab. If the properties are not correct, reconfigure the disk. Also, select the Policies tab and uncheck the box for "Enable write caching on the disk" so that this is disabled, as this is required by Oracle.
- Use Windows Disk Management to configure the shared disk storage partitions. Right-click the new disk from step 2 and select Properties. Create one extended partition of maximum size.
  - Right-click the partition and choose to Create Logical drive. Choose a minimum of the required size for the first logical drive partition, based on the configuration choices made previously. Do not choose a driver letter and do not format the drive. The drive letter will be assigned by Oracle during the installation of the Oracle software. Repeat this procedure for each of the remaining required logical drive partitions, if any.
  - On each of the other server nodes, use Windows Disk Management to verify that the shared storage partition(s) can be seen, as this is required before installing the Oracle Cluster Ready Services software in the next section. It may be necessary to reboot the nodes to see the same set of shared storage partitions from each of the other nodes. If any drive letters are seen, they must be removed, using the Windows Disk Management tool.

## Oracle Cluster Ready Services (CRS) software installation

1. Test the network communication between the nodes by using the ping utility for both the public and private names of the nodes (the names entered in the hosts file). Also verify administrative access to all nodes by executing the following commands from each node at a command prompt:

```
NET USE \\<host name>\c$
```

where <host name> is the public name of the other node. The command should complete successfully.

2. Create either a local image or a shared Oracle home for the Oracle 10g Database software from Node 1 of the cluster, by using following steps:
  - Obtain the Oracle 10g (10.1.0.2) Cluster Ready Services (CRS) software CD image from Oracle (1 zip file)
  - Unzip the CD image file on the Node 1 local drive, which creates a crs folder
  - Start the Oracle Universal Installer (OUI) by running the setup.exe file at the root of crs
  - Note: The CRS home cannot be installed on OCFS volumes, but must be on the local drive of each server node
  - Select the correct choices for disk formatting, based on previous configuration choices. This is where drive letter(s) are selected for OCFS
3. If using ASM for Oracle Database data files, then the logical drives must be stamped using the asmtoolg utility in the CRS home bin directory (see RAC Guide, chapter 7) at this point, before installing the 10g Database software.

## Oracle 10g Database Enterprise Edition software installation

1. Install the Oracle 10g Database Enterprise edition software from Node 1 of the cluster, by using following steps:
  - Obtain the Oracle 10g (10.1.0.2) Database software CD image from Oracle (1 zip file)
  - Unzip the CD image file on the Node 1 local drive, which creates a Disk1 folder
  - Start the Oracle Universal Installer (OUI) by running the setup.exe file at the root of Disk1
  - Select the correct choices for installation destination, depending on configuration choices made previously (e.g., drive letter for shared Oracle home, if used; otherwise, local drive letter that is the same on each of the nodes in the RAC cluster)
  - Check the box to “Do not create a starter database” so that only the software is installed and no sample database is created. An Oracle patchset will be applied to the installation to fix some issues with 10.1.0.2 before using Oracle’s DBCA utility to create a sample database.

## Post-installation configuration

### Apply Patchset 10.1.0.4 Patch Set 2 for 32-Bit Windows Server 2003

1. Obtain Oracle’s 9.2.0.4 patchset # 4163362 for 32-bit Windows from MetaLink, unzip the file, and copy it to a folder (e.g. c:\Patches) on the local drive on Node 1 (see PatchSet document).

2. Stop all existing Oracle processes on each node (if you do not do this at this time, you are prompted to do so during the OUI installation process).
3. Run the setup.exe utility to start the OUI Welcome screen.
4. Choose the destination as the Oracle home for CRS, so that the patchset files will be applied to the CRS software first.
5. After the successful installation of the Patchset to CRS files, repeat the process of starting the OUI. This time, select the destination as the Oracle home for the Database software (either shared or on each local server's drive) and apply the Patchset to the Database files.
6. After the patch set installation completes, click Exit and confirm it.

## Creation of Sample Oracle 10g Database

7. You must run the vipca.bat tool to configure virtual addresses before creating the sample database. If you do not do it now, you will be prompted to do so later. The tool is found in the Oracle 10g Database home (either on each server or on the Oracle shared home) in the bin directory. Run the vipca.bat utility for each Oracle home and follow the instructions.
8. On Node 1, run Oracle's Database Configuration Assistant (DBCA) to create a sample database to confirm a successful RAC installation. Run DBCA from the Start->All Programs->Oracle - OraDb10g - home1->Configuration and Migration Tools->Database Configuration Assistant. This starts the GUI screen. Select Create Database and all nodes (see RACGuide, chapter 11 for this section).
9. Select the General Purpose database type and assign a SID (e.g., ORCL).
10. Select Enterprise Manager Database control for database management, if desired.
11. Assign password(s) as desired.
12. Select database type—either OCFS or ASM, based on configuration choices made previously. Follow different paths of choices, depending on this choice.
13. During the DBCA creation process, an error may be displayed about a failure to configure Enterprise Manager and how to do this later, manually. Click ok to finish the database creation process.
14. After completing password management, the installation should complete successfully and start the cluster and all database instances. Another error may be seen indicating a failure to start the instance on a node other than the first node. A note on MetaLink (311321.1) suggests a recovery for this problem by running srvctl to alter a parameter for database startup. Also, rebooting the node(s) may fix the problem.
15. Confirm that the path environmental variable on each server node has been edited to include the location of the Oracle home\bin directory.
16. After all instances have been created on the nodes and are up, verify the database is running correctly in the cluster by using SQLPLUS at a command prompt with the following commands on each node:

```
Sqlplus /nolog
SQL> connect / as sysdba
SQL> select * from v$thread;
```

The output displayed in the command window should list all the nodes with a status of open, once they have all been started.

```
SQL> select * from v$instance;
```

The output should list a row for each node and give the names of all the nodes, after all the instances have been started.

This completes the installation of the Oracle 10g RAC database on Windows Server 2003.

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Java is a US trademark of Sun Microsystems, Inc.

Intel is a trademark of Intel Corporation in the U.S. and/or other countries.

Itanium is a trademark or registered trademark of Intel Corporation in the U.S. and other countries and is used under license.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corp.

Oracle is a registered US trademark of Oracle Corporation, Redwood City, California.

[10/2004]-1