



Next Generation automation

TODAY'S SOFTWARE COMPANY
TOMORROW'S METHODOLOGY AND TECHNOLOGY

Proof of Concept

Oracle® 10g R2 RAC on Windows Server® 2003 using Wasabi® Storage Builder® for iSCSI 1500i

© 2006 Next Generation automation

Roetaard 72
2547 Lint
Belgium

<http://www.nga.be>

Prepared by:	De Schryver Pierre [pschryve@nga.be]
Document Version:	1.0
Date:	06 March 2006
Revision Number	122



Version	Date
1.0	06 March 2006

Disclaimers

This document is the proprietary and exclusive property of Next Generation automation except as otherwise indicated. No part of this document, in whole or in part, may be reproduced, stored, transmitted, or used for design purposes without the prior written permission of Next Generation automation.

The information contained in this document is subject to change without notice.

The information in this document is for information purposes only. Next Generation automation disclaims all warranties, express or limited, including, but not limited, to the implied warranties of merchantability and fitness for a particular purpose, except as provided for in a separate software license agreement.

Trademarks

ORACLE® Oracle, JD Edwards, PeopleSoft, and Retek are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

Microsoft® Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.



Wasabi®, Flashware®, Storage Builder®, and Storage in a Flash® are registered trademarks of Wasabi Systems Inc.

Document History

Version	Issued By	Issue Date	Comments
1.0	De Schryver Pierre	06 March 2006	No previous document



Table of Contents

DOCUMENT HISTORY	2
0. GENERAL INFORMATION	4
1. ENVIRONMENT	5
2. CONFIGURATION BEFORE INSTALLATION	9
3. SUMMARY ORACLE INSTALLATION AND SETUP	13
4. TEST SCENARIOS	16
5. CONCLUSION	20



0. General Information

0.1. Background

This document provides a proof of concept (POC) for an Oracle Database 10g Release 2 Real Application Clusters (RAC) on Windows Server 2003 solution, using a Wasabi Storage Builder for iSCSI 1500i IP-SAN. It should be noted that even though this example uses iSCSI, the steps are almost identical for an implementation using a FC SAN for host storage connectivity. The POC is to be used as a template serving as a reference for creating a customer implementation.

Microsoft Cluster Software (MSCS) is not required for RAC databases as the Oracle Clusterware provides the clustering. However, the Oracle Clusterware can coexist with MSCS as long as the quorum and shared disks are partitioned and mutually exclusive.

During the Oracle installation some pitfalls occurred for which workarounds have been found and documented. For this, a detailed “**installation guidelines**” can be ordered.

References:

Oracle® Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide 10g Release 2 (10.2) for Microsoft Windows (Part Number B14207-04).

Oracle 10g Grid & Real Application Clusters – Oracle10g Grid Computing with RAC – Rampant TechPress (ISBN 0-9744355-4-6)

0.2. Overview of the Document

The document includes following sections:

- Introduction: Summary of the background and description of each section of the document.
- Environment: Overview of the RAC 10g topology and the specification of the software and hardware that will be used.
- Configuration before installation: Description of the configuration procedures before installing the Cluster Ready Services (CRS) and the RAC software.
- Summary Oracle installation and setup: Cluster Ready Services, Oracle database 10g software, Oracle database 10g companion CD software, Oracle network setup, Oracle database creation.
- Test Scenarios: Tablespace creation time and average write performance, TPC-B benchmark, Burn-in test.
- Conclusion.



1. Environment

1.1. 2-Node RAC 10g Topology

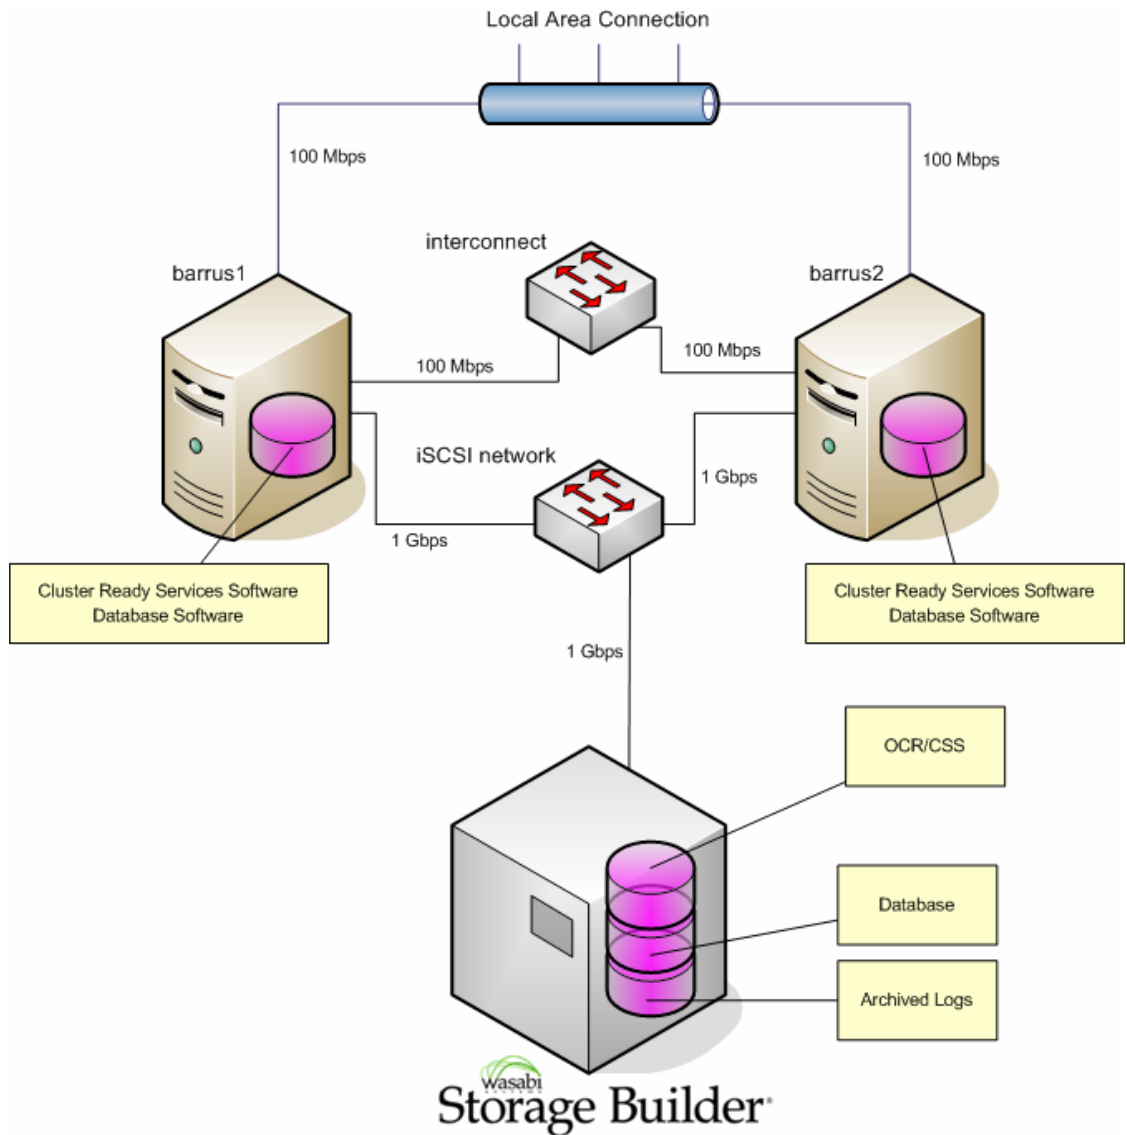


Figure 1- 2-Node RAC 10g Technology



1.2. Software specifications

1.2.1. System

Specification	Value
System architecture	32-bit
Operating system	Windows Server 2003 with service pack 1
Operating system language	English
Network protocol	TCP/IP

Table 1 – System Software specifications

1.2.2. Oracle

Specification	Version
Oracle Cluster Ready Services (Release 2)	10.2.0.1
Oracle Database 10g (Release 2)	10.2.0.1
Oracle Database 10g Companion CD (Release 2)	10.2.0.1

Table 2 – Oracle Software specifications



1.3. Hardware specifications

1.3.1. Nodes

Specification	barrus1	barrus2
Node	Supermicro P4SCE	Supermicro P4SCE
Processor	Pentium IV 3.2 Ghz	Pentium IV 3.2 Ghz
Local hard disk	2 x 120 GB	2 x 120 GB
Memory	1GB	1GB
Network adapter 1	3Com 3C905TX / 100 Mbps	3Com 3C905TX / 100 Mbps
Network adapter 2	Intel PRO/1000 MT Desktop Adapter / 1000 Mbps	Intel PRO/1000 MT Desktop Adapter / 1000 Mbps
Network adapter 3	Intel PRO/1000 MT Network Adapter / 1000 Mbps	Intel PRO/1000 MT Network Adapter / 1000 Mbps

Table 3 – Nodes specifications

1.3.2. Storage Area Network

Specification	Value
Model	Wasabi Storage Builder 1500i
Capacity Raid 1	236483 MB
Capacity Raid 50	946107 MB
Node	Supermicro 833T
Memory	1 GB
Processor	Xeon 2.8 Ghz
Hard-disks	8 x 250 GB SATA RAID 7200T

Table 4 - Storage Area Network Hardware

LUN	Destination	Size	RAID Level
1	OCR/CSS	300 MB	50
2	Database data files	25 GB	50
3	Archived logs	6 GB	50

Table 5 - Storage Area Network LUN's



1.3.3. Network Switches

Specification	Model	Speed
Private Interconnect	D-Link DGS-1008D	10/100/1000 Mbps
Public LAN	D-Link DES-1016R	10/100 Mbps
iSCSI	D-Link DGS-1216T	10/100/1000 Mbps

[Table 6 – Network Switches](#)



2. Configuration before installation

2.1. Network

2.1.1. Network Interface names

Name	Purpose
LAN	Public network
INTERCONNECT	Private interconnect between barrus1 and barrus2
iSCSI	iSCSI connection IP-SAN

Table 7 - Network Interfaces names

2.1.2. Public network specification

Host Name	IP address	Netmask	Type
barrus1	10.72.1.220	255.255.0.0	Public
barrus2	10.72.1.221	255.255.0.0	Public
barrus1-vip	10.72.1.222	255.255.0.0	Virtual
barrus2-vip	10.72.1.223	255.255.0.0	Virtual

Table 8 - Public network specification

2.1.3. Private network specification

Host Name	IP address	Netmask	Type
barrus1-priv	10.1.1.220	255.255.255.0	Private
barrus2-priv	10.1.1.221	255.255.255.0	Private

Table 9 - Private network specification



2.1.4. iSCSI network specification

Host Name	IP address	Netmask	Type
barrus1	10.10.2.220	255.255.255.0	iSCSI
barrus2	10.10.2.221	255.255.255.0	iSCSI

Table 10 - iSCSI network specification

2.2. Disk System

2.2.1. Overview created LUN's

ID	Role	Description (Alias)	REMOVE
0	Target	San-1-01	REMOVE
1	Target	San-1-02	REMOVE
2	Target	San-0-01	REMOVE
3	Target	OCR/CSS	REMOVE
4	Target	Oracle Datafiles	REMOVE
5	Target	Oracle Archived Logs	REMOVE

Figure 2 - Overview created LUN's



2.2.2. SAN disks

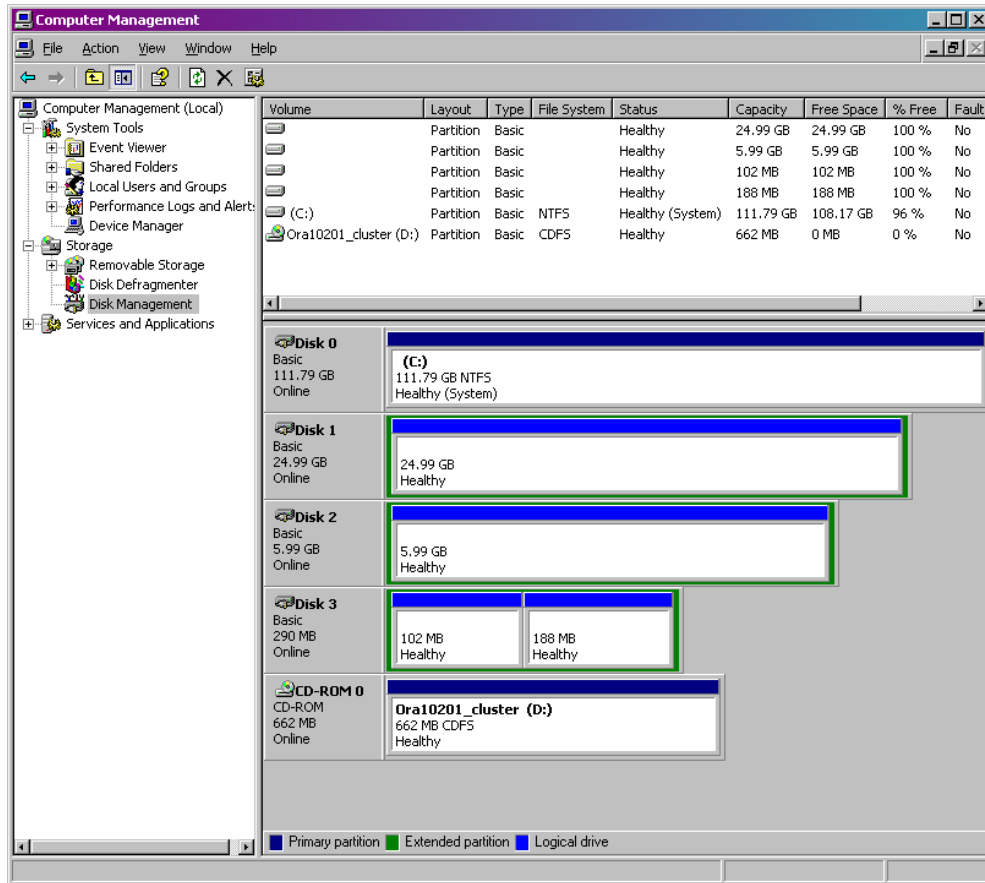


Figure 3 - Disk Management SAN Disks

Cluster Ready Services

LUN	Partition Type	Size	Logical Drive	Size
1	Extended	300 MB	First	100 MB
1	Extended		Second	200 MB

Table 11 - Cluster Ready Services SAN Disks



Database storage

LUN	Partition Type	Size	Logical Drive	Size
2	Extended	25 GB	First	25 GB
3	Extended	6 GB	First	6 GB

Table 12 - Database storage SAN Disks

2.2.3. Local Oracle Software and administration Disks

Software	Path	Size
Cluster Ready Services software	C:\oracle\product\10.2.0\crs	500 MB
Database software	C:\oracle\product\10.2.0\db_1	1 GB
Administration	C:\oracle\product\10.2.0\admin	500 MB

Table 13 - Oracle Software Local Disks



3. Summary Oracle installation and setup

3.1. Cluster Ready Services

Following Oracle services should run on both nodes after successful install and setup.

Oracle Services
Oracle Object Service
OracleClusterVolumeService
OracleCRService
OracleCSService
OracleEVMService

Table 14 - Oracle Services

3.2. Oracle Database 10g Software

Custom installation type

Components	Select Status
Data Mining Scoring Engine 10.2.0.1.0	✓
Oracle Windows Interfaces 10.2.0.1.0	✓
Oracle Services for Microsoft Transaction Server 10.2.0.1.0	✗
iSQL*Plus 10.2.0.1.0	✓

Table 15 - Components Custom Installation mode



3.3. Database 10g Companion CD Software

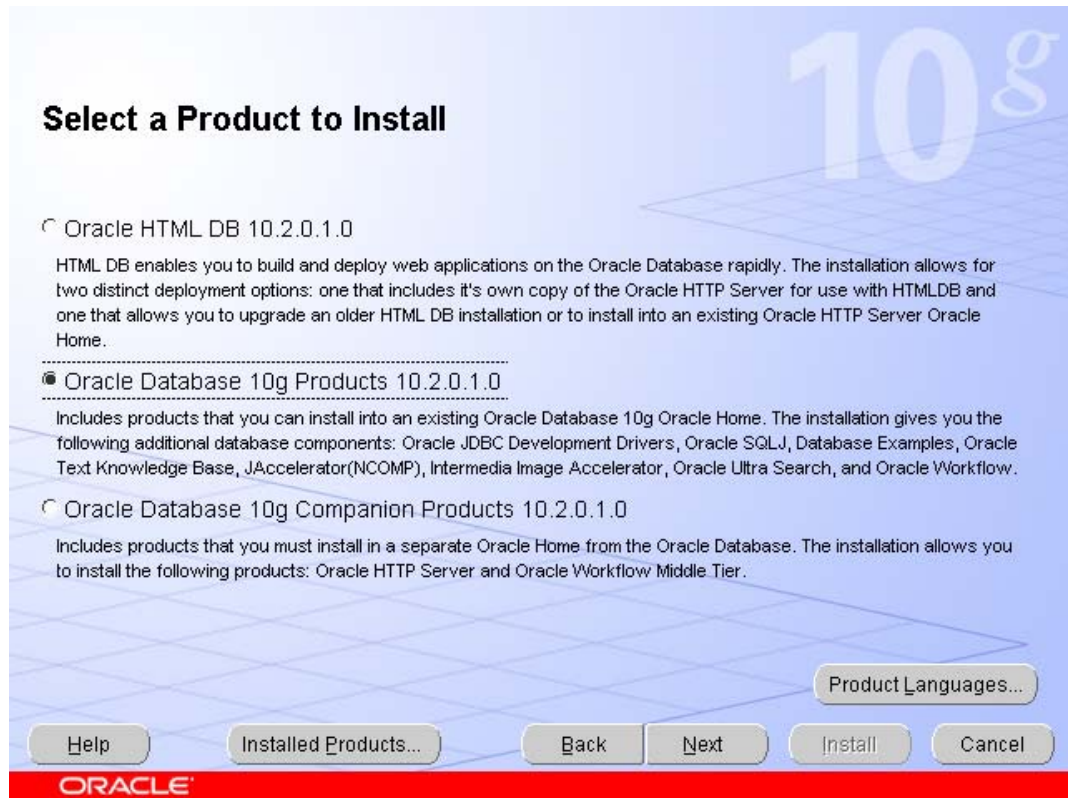


Figure 4 - Product to Install from the Companion CD

3.4. Oracle network setup

Following Oracle listener services should run after successful setup.

Node	Service
barrus1	OraDb10g_home1TNSListenerLISTENER_BARRUS1
barrus2	OraDb10g_home1TNSListenerLISTENER_BARRUS2

Table 16- Oracle listener services



3.5. Oracle database

→ Database NGA09T (Instance NGA09T1) created.

Object type	Object/File Name	Size
Tablespace	SYSTEM	300 MB
	SYSAUX	120 MB
	TEMP	20 MB
	UNDOTBS1	200 MB
	UNDOTBS1	200 MB
	USERS	5 MB
	EXAMPLE	150 MB
	Redo logs	R:\NGA09T\redo01_01_1.log
R:\NGA09T\redo02_01_1.log		250 MB
R:\NGA09T\redo03_01_2.log		250 MB
R:\NGA09T\redo04_01_2.log		250 MB
R:\NGA09T\redo05_01_1.log		250 MB
R:\NGA09T\redo06_01_2.log		250 MB

Table 17 - Database specifications

- Archiving activated.
- Second Database Instance (NGA09T2) activated.



4. Test Scenarios

In the following subsections some tests have been executed upon the successful build of the solution described in this document (NGA09T) and upon one-instance database running entirely on local disks (NGA10T).



These tests are only for information and should not be interpreted as an official benchmark statement. We are aware that better performance can be achieved after proper tuning.

4.1. Tablespace creation time and average write performance

Tablespace	Size	NGA09T		NGA10T	
		Time	Mb/sec	Time	Mb/sec
SINGLE_DATAFILE	1 data file of 20 GB	4m:57s	69	6m:34s	52
MULTI_DATAFILES	10 data files of 2 GB	4m:37s	74	6m:22s	54

Table 18 - Tablespace creation



4.2. TPC-B benchmark

The TPC-B benchmark measures transaction in terms of how many transactions per second (TPS) a system can handle. The TPC-B benchmark can be compared to an electronic data processing batch application that runs overnight when no customer users are logged in.

Counter	Value
User loads	10 users increment up to 100 users
Pre-sampling time	1m:0s for each new increment
Sampling time	3m:0s for each new increment

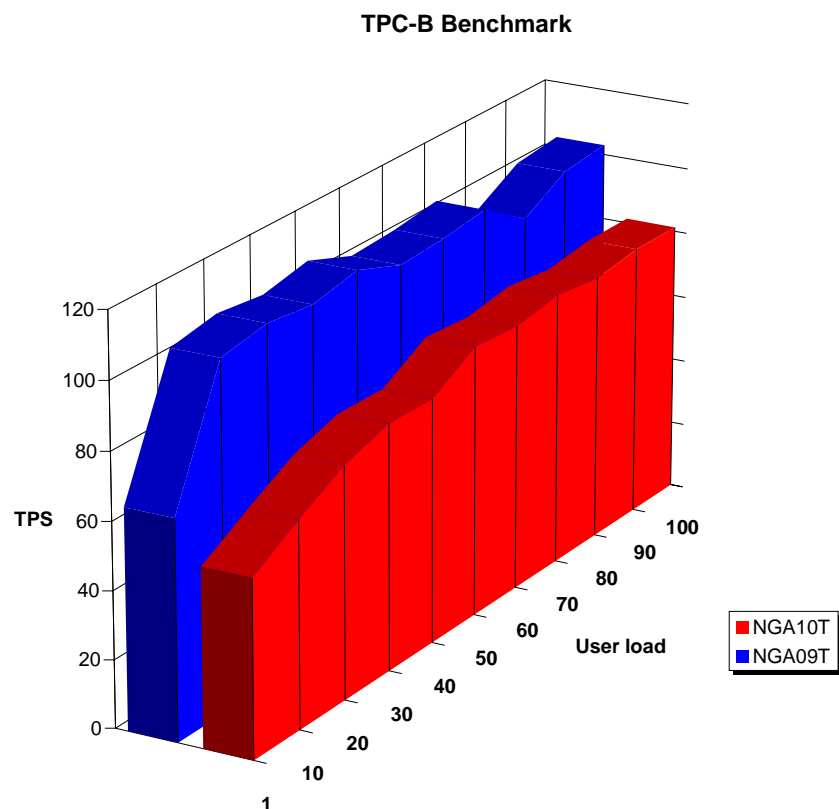


Figure 5 - TPC-B benchmark



4.2.1. NGA09T (RAC)

User load	TPS	Avg. Time	Avg. Read Bytes/Sec		Avg. Write Bytes/Sec	
			NGA09T1	NGA09T2	NGA09T1	NGA09T2
1	64.82	0.015	1072910	45853	693711	17158
10	102.50	0.097	845342	668907	541172	504790
20	105.10	0.189	854627	589165	577191	522576
30	103.20	0.289	599978	597171	573688	561487
40	106.10	0.375	700282	576512	577324	518257
50	100.60	0.495	718963	572247	525091	558154
60	101.40	0.589	670700	548070	550932	554737
70	103.20	0.675	567550	551422	541736	516715
80	93.68	0.851	773875	615466	524735	521518
90	101.30	0.884	715596	541277	533347	551100
100	102.90	0.968	567483	554963	559146	518729

Table 19 - TPC-B Benchmark on NGA09T

4.2.2. NGA10T (Single Instance)

User load	TPS	Avg. Time	Avg. Read Bytes/Sec	Avg. Write Bytes/Sec
1	52.51	0.018	711071	719979
10	60.89	0.163	712976	718394
20	68.69	0.289	773152	728243
30	73.07	0.407	806792	759103
40	72.65	0.547	783203	742847
50	80.50	0.615	822411	824465
60	79.32	0.750	816275	830884
70	81.46	0.852	825001	840267
80	79.50	0.997	796657	817342
90	81.38	1.096	821262	843714
100	81.26	1.219	818877	844067

Table 20 - TPC-B Benchmark on NGA10T



4.3. Burn-in test

Burn-in tests are used to test databases for endurance and reliability with a given load over several hours (or even days).

Counter	Value	
User loads	200	
Pre-sampling time	4m:0s	
Sampling time	12h:0m:0s	
Test type	Mixed Workload Database Test	
Avg. TPS	68.94	
Total executions	2978092	
Deadlocks	0	☹️
Rollbacks	0	☹️
Errors	0	☹️

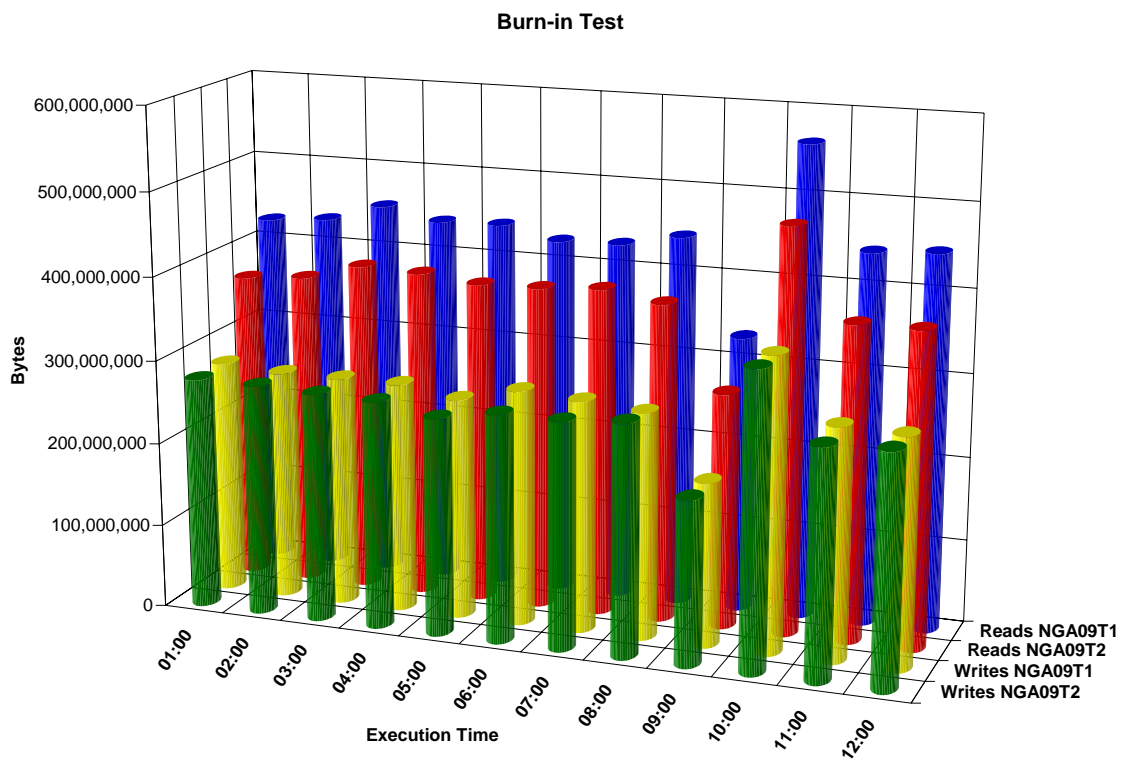


Figure 6 - Disk IO Bytes/Hours



5. Conclusion

Running Oracle Database 10g Release 2 Real Application Clusters (RAC) on Windows Server 2003 solution, using a Wasabi Storage Builder for iSCSI 1500i IP-SAN is **feasible**.

Nevertheless some remarks should be taken into consideration:

- During the Oracle installation some pitfalls occurred for which workarounds have been found and documented.
- As stated before, this is a POC and not a benchmark. Results will surely improve after database tuning and using appropriate hardware like an iSCSI HBA instead of a network adapter.
- Oracle Note 220970.1: *For iSCSI, Oracle has made the statement that, as a block protocol, this technology does not require validation for single instance database. There are many early adopter customers of iSCSI running Oracle9i and Oracle Database 10g. As for RAC, Oracle has chosen to validate the iSCSI technology (not each vendor's targets) for the 10g platforms - this has been completed for Linux, Unix and Windows. For Windows we have tested up to 4 nodes - Any Windows iSCSI products that are supported by the host and storage device are supported by Oracle. No vendor-specific information will be posted on Certify. -- Ref #: ID-5788 --*