



ORACLE®

Best Practices for Oracle Database on Windows

Karin Brandauer, Principal Support Engineer
Bug Diagnostics and Escalations, Oracle Global Support



Agenda

- Oracle Architecture on Windows
- Increasing Addressable Process Memory
- Tuning Memory with AWE/PAE, ORASTACK and Shared Server
- Monitoring Memory Usage
- 64-bit Environments
- General Tips on Running Oracle/Windows
- High Availability Solutions for Scalability
- Maximum Availability Architecture (MAA)
- Patching



Windows 32-bit Platform Support

OS	9iR2	10gR1	10gR2
Windows 2000	Yes	Yes	Yes
Windows XP Professional	Yes	Yes	Yes
Windows Server 2003	Yes	Yes	Yes
Windows Vista	No	No	Yes*

*Requires special 10203 client/server release for Vista



Windows 64-bit Platform Support

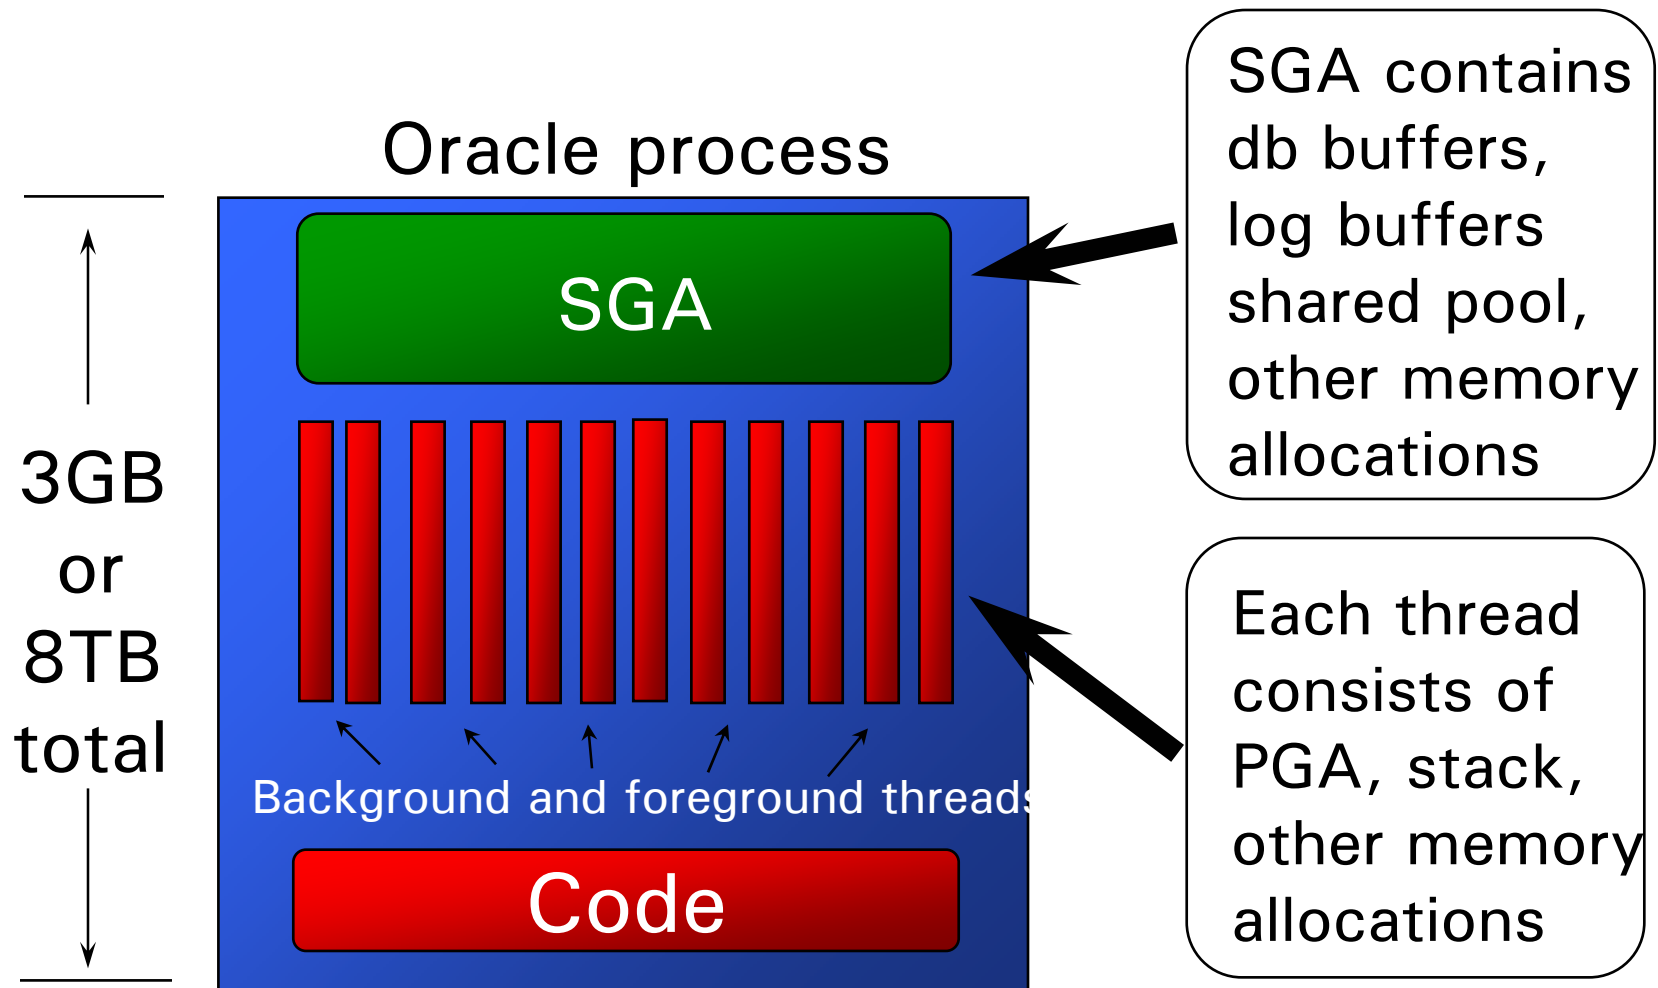
OS	9iR2	10gR1	10gR2
Windows Server 2003 for Itanium (64-bit DB)	Yes	Yes	Yes
32-bit Windows XP & Windows Server 2003 on EM64T (32-bit DB)	Yes	Yes	Yes
Windows XP & Windows Server 2003 x64 (32-bit DB)	Yes (client only)	Yes (client only)	Yes (client only)
Windows XP & Windows Server 2003 x64 (64-bit DB)		Developer Release (May 04)	Yes
Windows Vista x64			Planned later CY2007




Database Architecture

- Unlike UNIX ports, Oracle on Windows is implemented as a single operating system process.
- Typical “processes” such as PMON and LGWR have been converted to native Windows threads running in a single process.
- One process exists per instance on a Windows server.

Architecture: Thread Model



32-Bit Address Space



Virtual Memory Address Space is limited to 4 GB in 32-bit architecture



4GT RAM Tuning

- Increase addressable memory available to the Oracle process by adding /3GB switch to boot.ini file:

```
multi(0)disk(0)rdisk(0)partition(1)\WINNT="Microsoft Windows 2000  
Advanced Server" /fastdetect /3GB
```

- Reboot server to enable
- Must monitor kernel memory closely to prevent instability of operating system
- See Metalink Notes 46001.1 and 297498.1
- See Microsoft KB article 297812



Monitoring Memory

- Key Items to Monitor for Memory Usage:
 - Performance Monitor - Virtual Bytes for oracle.exe to see total memory used by the process
 - Total Pool Non-Paged Bytes – Memory Counter
 - Total Pool Paged Bytes – Memory Counter
 - Free System Page Table Entries (PTE's) – Memory Counter
 - In addition to Perfmon, the Process Explorer tool from Windows Sysinternals is invaluable:

<http://www.microsoft.com/technet/sysinternals/default.mspx>

Performance Monitor

The screenshot shows the Windows Performance Monitor application. The 'Add Counters' dialog box is open, allowing the user to select performance counters. The dialog is configured to select counters from the local computer (\server1) for the 'Process' performance object. The 'Select counters from list' option is chosen, and the 'Virtual Bytes' counter is selected from the list. The 'oracle' instance is selected from the list of instances. The background shows a graph with a vertical red line at approximately 47% and a table with columns for Color, Scale, Counter, and Instance.

Color	Scale	Counter	Instance
—	0.000...	Virtual By...	oracle

Process Explorer

Process Explorer - Sysinternals: www.sysinternals.com

File Options View Process Find Handle Help

Process	PID	CPU	Description	Company Name	Path
WinMgmt.exe	1016		Windows Management Instrumentation	Microsoft Corporation	C:\WINNT\system32\wbem\WinMgmt.exe
nutsrv4.exe	1092		NuTcracker Service	DataFocus, Inc.	C:\WINNT\system32\nutsrv4.exe
telnetd.exe	1212		telnetd service	DataFocus, Inc.	C:\WINNT\system32\telnetd.exe
svchost.exe	1304		Generic Host Process for Win32 Services	Microsoft Corporation	C:\WINNT\system32\svchost.exe
cvpnd.exe	1412		Cisco Systems VPN Client	Cisco Systems, Inc.	C:\Program Files\Cisco Systems\VPN Client\cvpnd.exe
ORACLE.EXE	2356	0.96	Oracle RDBMS Kernel Executable	Oracle Corporation	d:\oracle\ora92\bin\oracle.exe
lsass.exe	264		LSA Executable and Server DLL (Export...)	Microsoft Corporation	C:\WINNT\system32\LSASS.EXE
Explorer.EXE	1252		Windows Explorer	Microsoft Corporation	C:\WINNT\explorer.exe
twatdog.exe	932				C:\WINNT\system32\TwaDog.exe
itouch.exe	1632		iTouch Application	Logitech Inc.	C:\Program Files\Logitech\iTouch\iTouch.exe
iclient.exe	1648	1.92	Integrity Client	Zone Labs Inc.	C:\Program Files\Zone Labs\Integrity Client\iclient.exe
qttask.exe	1664			Apple Computer, Inc.	C:\Program Files\QuickTime\qttask.exe
ccApp.exe	996		Symantec User Session	Symantec Corporation	C:\Program Files\Common Files\Symantec Shared\ccApp.exe
VPTray.exe	1240		Symantec AntiVirus	Symantec Corporation	C:\Program Files\Symantec AntiVirus\WPTray.exe
jusched.exe	1716		Java(TM) 2 Platform Standard Edition bi...	Sun Microsystems, Inc.	C:\Program Files\Java\jre1.5.0_04\bin\jusched.exe
ctfmon.exe	1732		Cicero Loader	Microsoft Corporation	C:\WINNT\system32\CTFMON.EXE
Hotsync.exe	848		HotSync® Manager Application	PalmSource, Inc.	C:\Program Files\Handspring\Hotsync.exe
DvzIncMsr.exe	1836		DataViz Update Checker	DataViz, Inc.	C:\Program Files\Common Files\DataViz\DvzIncMsr.exe
trillian.exe	512		Trillian	Cerulean Studios	C:\Program Files\Trillian\trillian.exe
thunderbird.exe	1764		Mozilla Thunderbird	Mozilla.org	C:\Program Files\Mozilla Thunderbird\thunderbird.exe
firefox.exe	1064		Firefox	Mozilla	C:\Program Files\Mozilla Firefox\firefox.exe
Palm.exe	1024		Palm Desktop Application	PalmSource, Inc.	C:\Program Files\Handspring\Palm.exe
Communicator.exe	556		Cisco IP Communicator	Cisco Systems, Inc.	C:\Program Files\Cisco Systems\Cisco IP Communicator\Communicator.exe
POWERPNT.EXE	1696		Microsoft Office PowerPoint	Microsoft Corporation	C:\Program Files\Microsoft Office\OFFICE11\POWERPNT.EXE

Type Name

- Directory \KnownDlls
- Directory \Windows
- Directory \BaseNamedObjects
- Event \BaseNamedObjects\userenv: User Profile setup event
- Event \BaseNamedObjects\crypt32LogoffEvent
- File D:\oracle\oradata\v92\UNDOTBS01.DBF
- File \Device\Tcp
- File \Device\Tcp
- File \Device\Np
- File \Device\Np
- File \Device\Np
- File \Device\NamedPipe\EVENTLOG
- File \Device\Tcp
- File D:\oracle\oradata\v92\DDM01.DBF
- File D:\oracle\oradata\v92\SYSTEM01.DBF
- File D:\oracle\ora92\rdbms\mesg\oraus.msb
- File D:\oracle\ora92\rdbms\mesg\oraus.msb
- File D:\oracle\ora92\rdbms\mesg\oraus.msb
- File D:\oracle\ora92\rdbms\mesg\oraus.msb

CPU Usage: 6.73% Commit Charge: 32.65% Processes: 51




OS Tools

- tasklist, taskkill
- tlist (Shows command line args with -c)
- driverquery
- diskpart
- sc (sc query state= all)

<http://www.microsoft.com/technet/sysinternals/default.mspx>

- regmon, filemon, procexp, tcpview
- Windows Services for Unix

32-bit: VLM Support



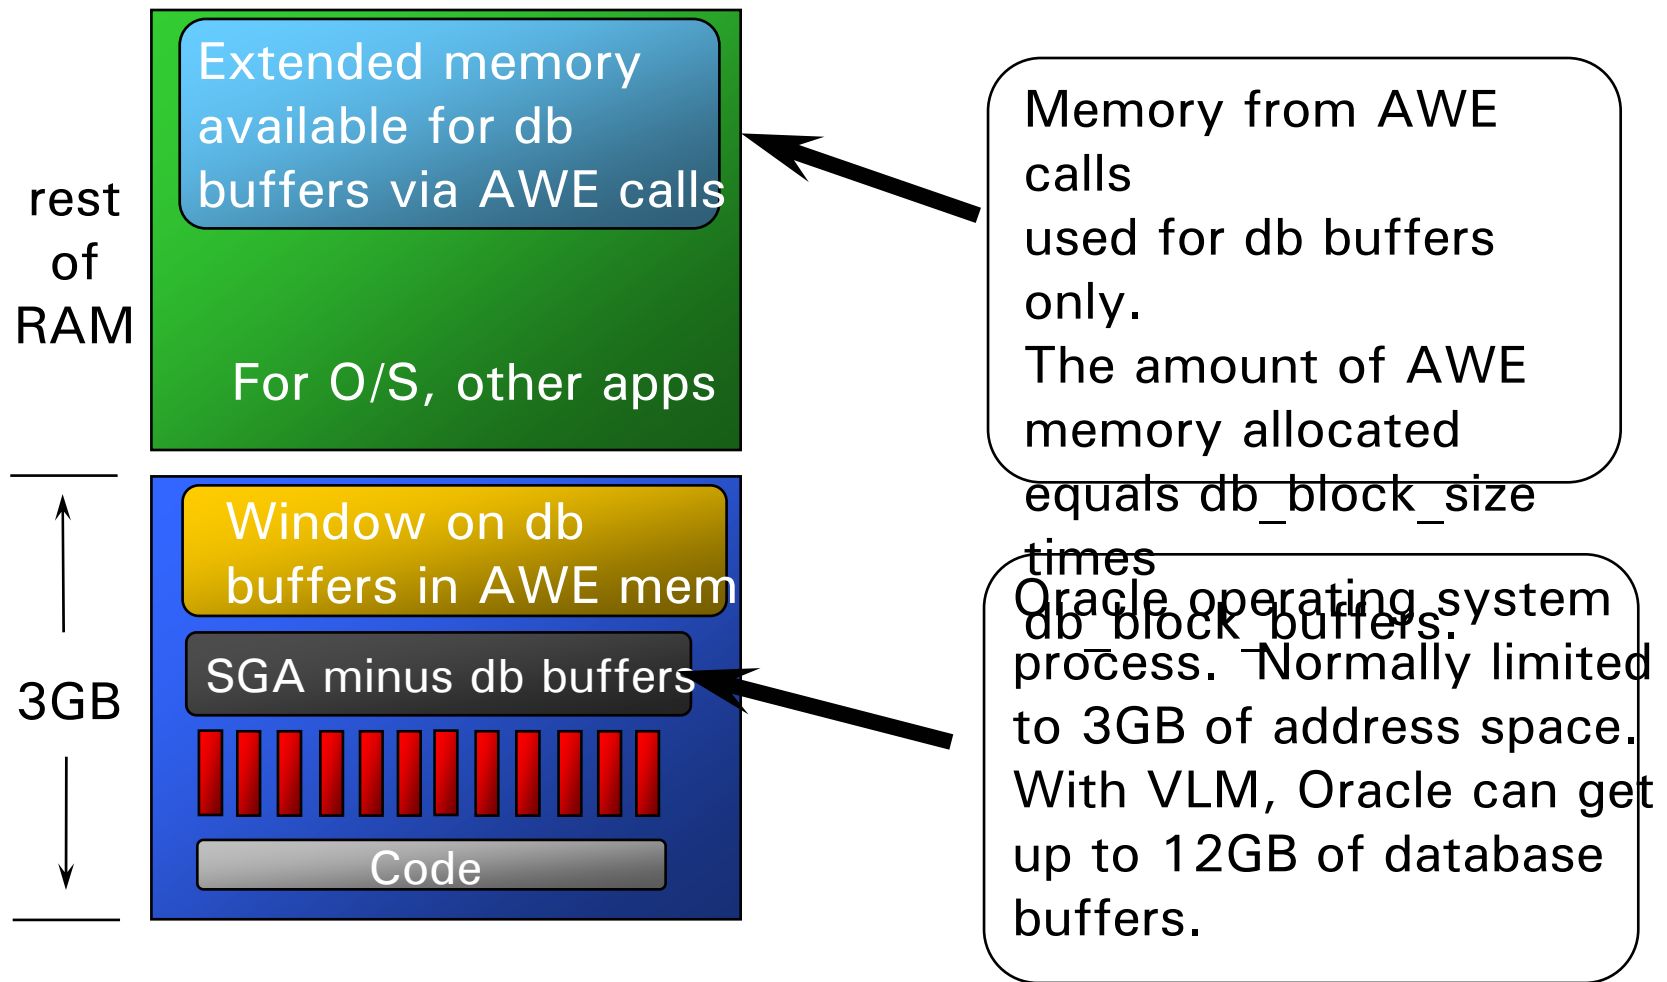
Windows Server 2003 Memory Limits (32-bit)

Standard Edition:
4GB

Enterprise Edition:
32GB

Datacenter Edition:
64GB

32-bit: VLM Support





Implementing AWE

- Use AWE with Oracle by adding initialization parameter `USE_INDIRECT_DATA_BUFFERS`
- Use `DB_BLOCK_BUFFERS` instead of `DB_CACHE_SIZE`
- With AWE, database buffer cache can be increased up to roughly 12 GB
- Default value for `AWE_WINDOW_MEMORY` is 1 GB
- See Metalink Note 225349.1 for more information



Using ORASTACK

- Each thread within Oracle process is provided 1MB reserved stack space
- Reduce to 500k without consequence:
C:\orastack tnslnr.exe 500000
C:\orastack oracle.exe 500000
- See Metalink Note 46001.1 for more information



Using Shared Server

- With Shared Server, client connections are ultimately sent to a Dispatcher, which binds the client to an already established virtual circuit
- Implement Shared Server in tnsnames.ora:
(DESCRIPTION=
 (ADDRESS=(PROTOCOL=tcp)
 (HOST=sales-server)(PORT=1521))
 (CONNECT_DATA= (SERVICE_NAME=sales.us.acme.com)
 (SERVER=shared)))
- See Net Admin Guide for more details



Best Practices for 32-Bit Memory Optimization

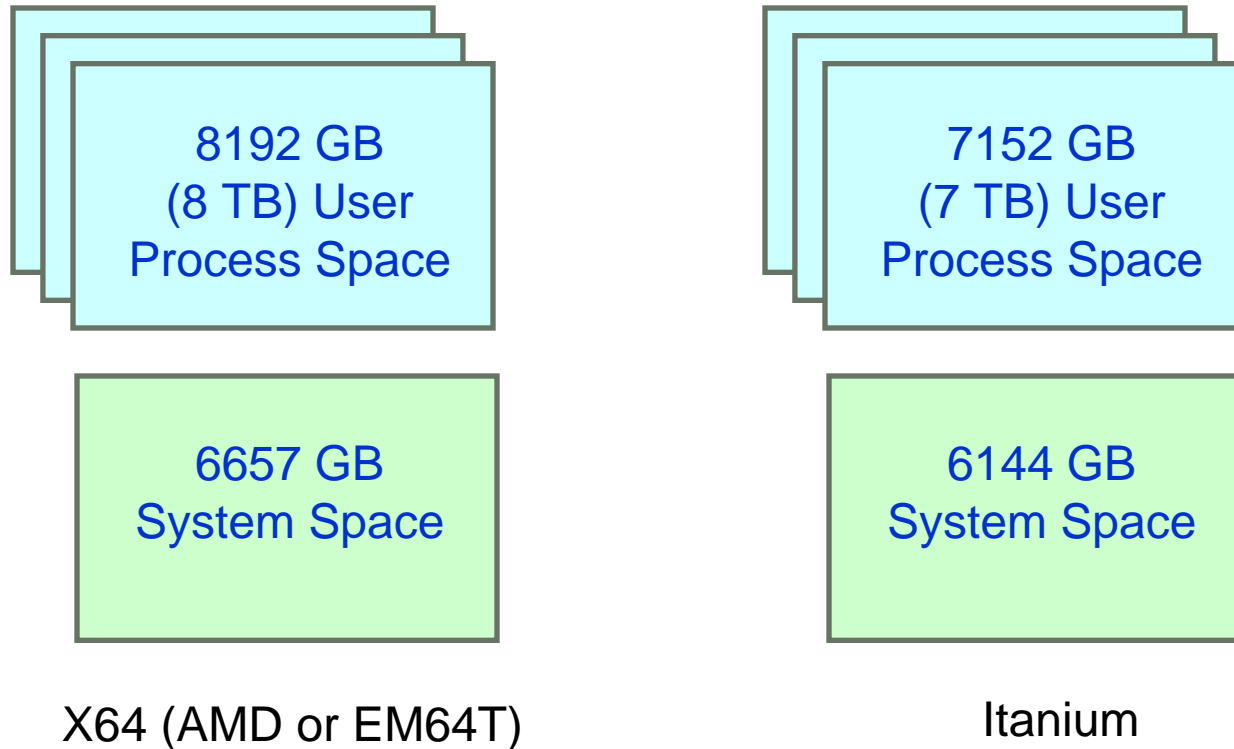
- Implement the /3GB switch
- Combine /3GB with /USERVA switch
- Run orastack to reduce stack size for both tnslsnr.exe and oracle.exe – be aware that patches change the executables, so run it again after patching.
- Use shared servers in implementations where large#’s of users connect in to the database.
- Control PGA Memory by using PGA_AGGREGATE_TARGET parameter



Best Practices for 32-bit Memory Optimization (continued)

- Verify Monitor Kernel Memory and Oracle Memory by using Perfmon or other tools that accurately measure Virtual Bytes.
- Use Automatic Workload Repository (AWR) to monitor cache hit ratios and shared_pool stats, etc. Make sure that values are not overstated
- If large buffer cache is needed, implement AWE, but be aware that using AWE disables Automatic Memory Management features (SGA_TARGET cannot be used when USE_INDIRECT_DATA_BUFFERS is set).

64-Bit Address Space



Virtual Memory Address Space is limited to 7-8 TB in 64-bit architecture (depending on chip)



64-bit Oracle Releases

- Oracle has a long history of supporting 64-bit databases on other platforms
- Interoperability between 32-bit clients and 64-bit servers and vice versa
- Improved performance, availability and scalability
- Itanium is supported with 9.2.0.3 and higher
- x64 (AMD/EM64T) is supported with 10.2.0.1 and higher



Migration to 64-bit

- 32-bit to 64-bit upgrade process is simple
 - 32-bit data files are compatible with 64-bit DB
- No need to recreate the database
- Full export and import not required
- Database Upgrade Assistant automates process
- Transparent migration for end-user applications
 - No changes required to existing client applications when running against 64-bit database



64-Bit Best Practices

- Run correct 64-Bit version of Oracle for the architecture – i.e. 64-Bit Oracle for AMD or 64-Bit Oracle for Itanium.
- 32-Bit Oracle RDBMS not supported on 64-Bit platforms
- Use MBR disks instead of GPT disks
- For RAC environments with SGA's > 4GB, apply latest Oracle patchset.
- Enable Large Pages (Note 422844.1)



Additional General Best Practices

- Take Advantage of Hyperthreading, which allows a single CPU to look like 2 CPU's
- Don't set ORACLE_HOME in environment – Oracle gets the environment via the Registry. Oracle.key file in OH\bin points to correct registry key.
- SQLNET.AUTHENTICATION_SERVICES=(NTS)
this is a default value in sqlnet.ora for a reason – it should be left at default.
- In 10.2.x, SQLNET.INBOUND_CONNECT_TIMEOUT defaults to 60 – this may need to be set to 0 in some situations. Prior to 10.2 this is not an issue. (Note 363705.1 explains)
- Use Automatic Storage Management (ASM) whether running single-instance or RAC



High Availability Solutions


- HA becomes essential as databases are critical component of business
- HA Goals: Minimize downtime to your company and your customers
- Solutions for Windows Environments
 - Oracle Fail Safe
 - Real Application Clusters (RAC)
 - Data Guard (DG)
 - Maximum Availability Architecture (MAA)



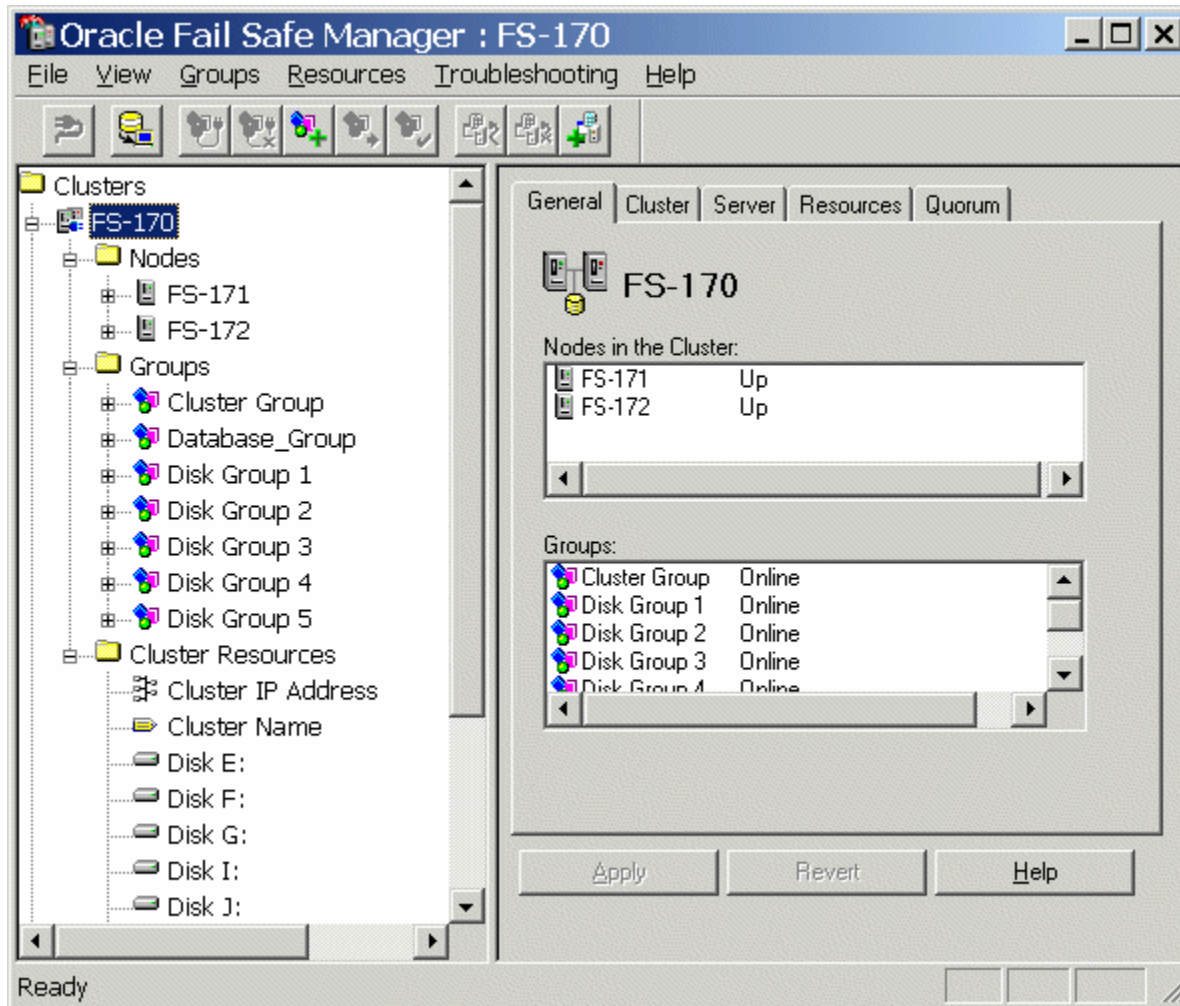
Oracle Fail Safe

- Integrated with Microsoft Clustering, Fail Safe is a core feature included with every Oracle 10g and Oracle9i license for Windows NT, Windows 2000, and Windows 2003
- In the event of a system failure, Oracle Fail Safe works with Microsoft Cluster Server to restart Oracle databases and applications on a surviving cluster node
- MSCS and Fail Safe uses “share-nothing” architecture (only one node can access shared datafiles at any time)

Fail Safe Architecture



Fail Safe Manager





Fail Safe Best Practices


- One database per group
 - Separate production from non-production databases into different groups
- Multiple physical disks to be separated into different groups
- Failback and Restart properties should be reviewed for business needs

<http://otn.oracle.com/tech/windows/failsafe/index.html>

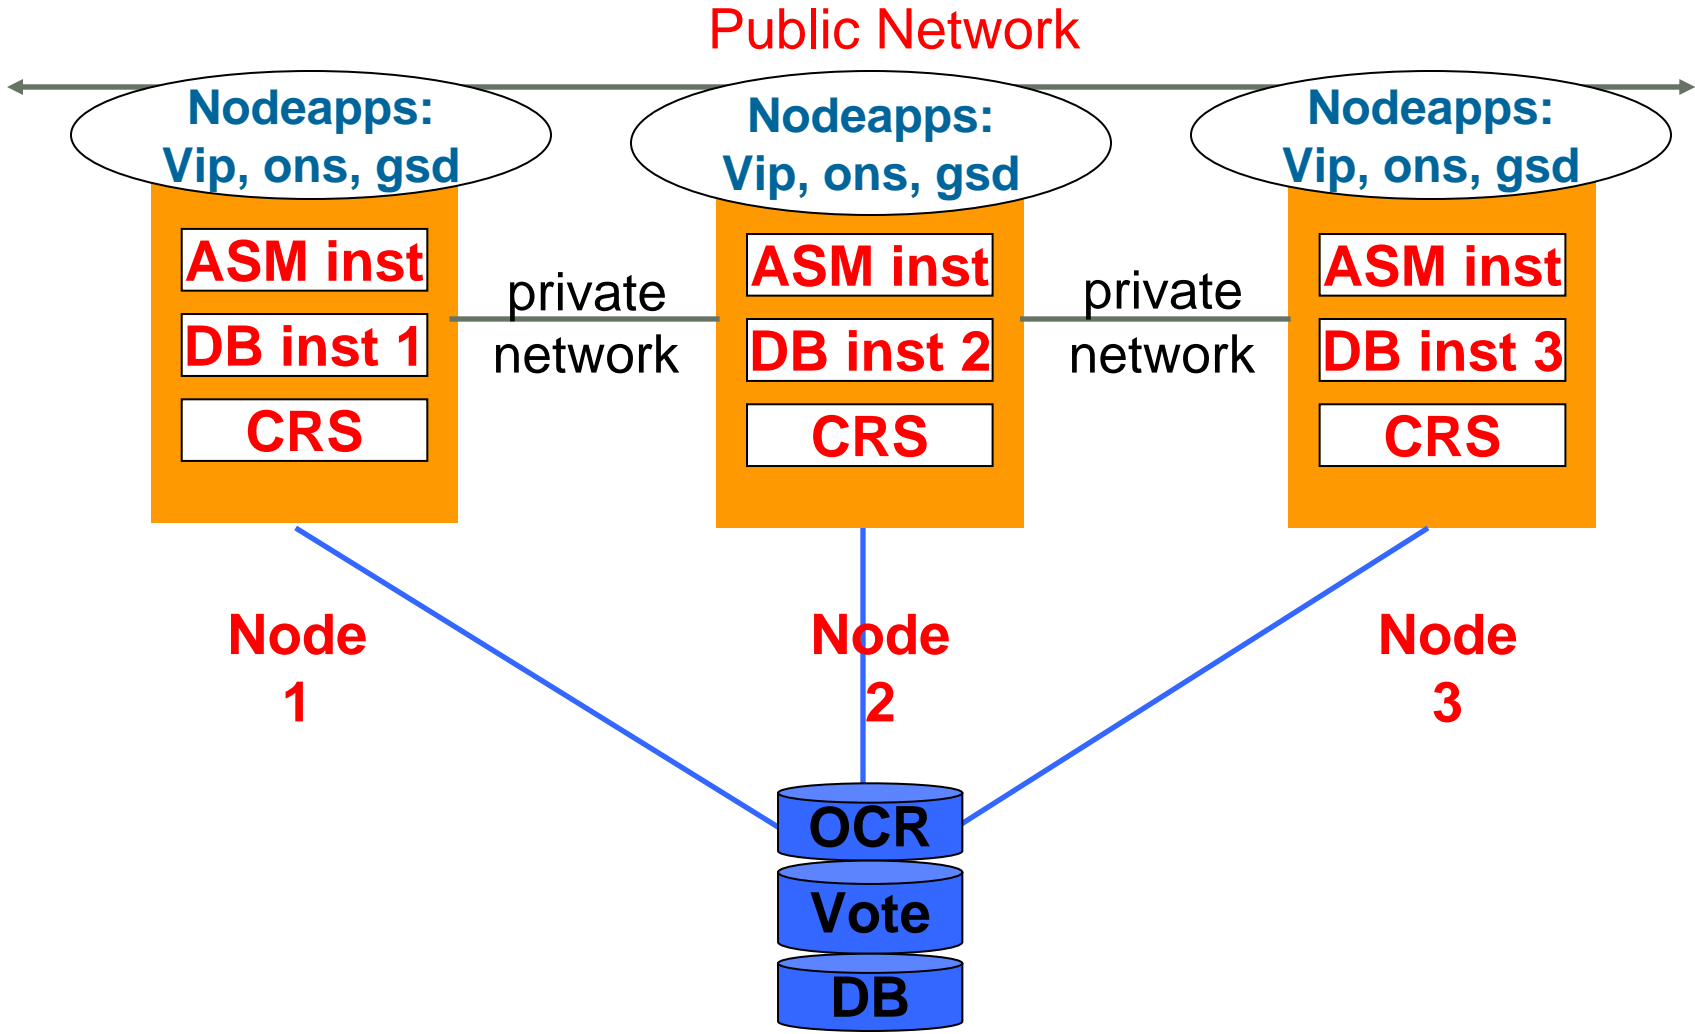


Real Applications Clusters

- Use RAC for scalability
 - Add instances against same database files providing more Oracle processes and increasing number of users
 - Provides unique scalability on Windows that no other vendor offers
- Clustered databases supported on Windows platforms since version 7.3.3
- Uses Oracle's own clustering software, not MSCS
- Oracle 10g provides platform independent Cluster Ready Services (CRS) to handle failover of services to surviving nodes



RAC Instance Architecture





RAC/CRS Install Options

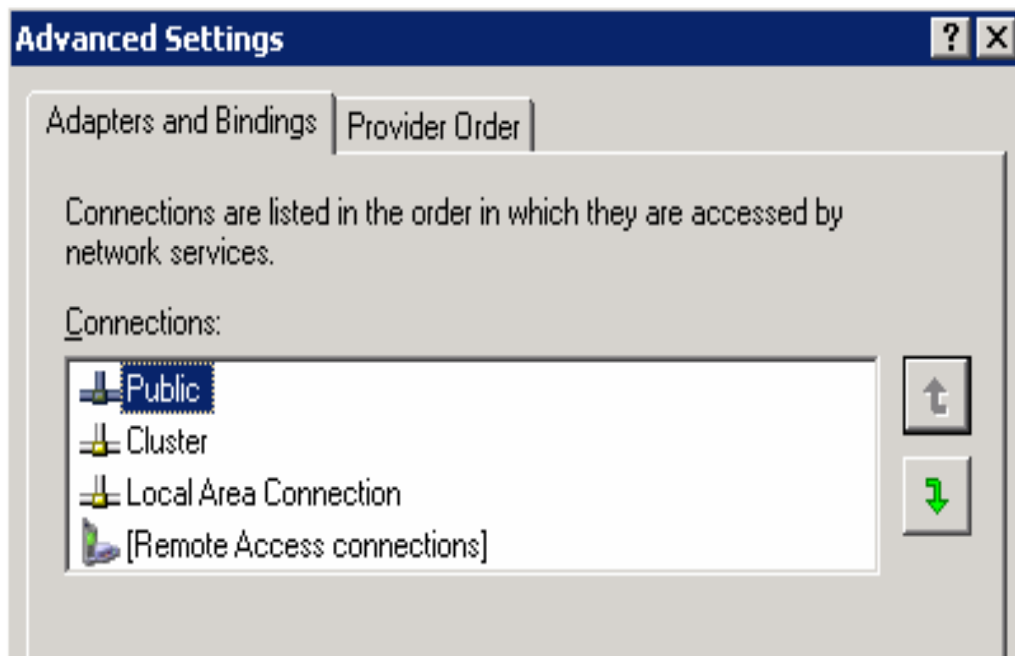
- Oracle Home can be on local NTFS drives or Oracle Cluster File System (OCFS)
 - Vast majority of installs use local NTFS homes
 - CRS home must be local NTFS drive
- Datafiles can reside on OCFS, RAW, or Automatic Storage Management (ASM)
- OCR and Voting file can reside on OCFS or RAW
- Refer to documentation for recommended partition sizes



Shared Storage

- Automount must be enabled in Windows 2003
 - Diskpart.exe can be used:
diskpart> automount enable
- Creating Partitions
 - Create all partitions on one node
 - Extended Partitions, w/logical drives
 - Use Basic disks: Dynamic Disks not supported
 - Windows 2003 requires a reboot
 - Remove drive letter assignments
 - Verify all nodes see all shared partitions

RAC Network Checklist



- Public adapter should be first
- Private adapter should be second
- Ping node's public hostname to verify
- Ping each node's public & private hostname

Network Connections

Advanced-> Advanced Setting



RAC Best Practices

- Eliminate Single Points of Failure
 - NIC's, Switches, Interconnect, Shared Storage, Power Supplies
 - Understand cost vs. availability tradeoff
- Use fastest switch available for private interconnect
 - Disable additional protocols such as spanning tree protocol
 - Increase MTU size as high as switch allows – i.e. 9000
 - Allow cards and switch ports to autonegotiate speed.
- Use static IP addresses
 - Public LAN resolved by DNS and hosts file
 - For cluster interconnect use non-routable IP (10.X or 192.168.X)



RAC Best Practices

- On 10gR1, be sure to manually backup Voting Disk/File.
- Backups of files on RAW devices can be done via 'ocopy' utility on windows – i.e.:
 - Ocopy \\.\votedsk C:\backups\votedsk.bak
 - Backups of OCR can be taken manually as well, but CRS backs up OCR every 4 hours (on 10.1.0.4)
- 10gR2 allows mirroring of OCR and Voting Disks

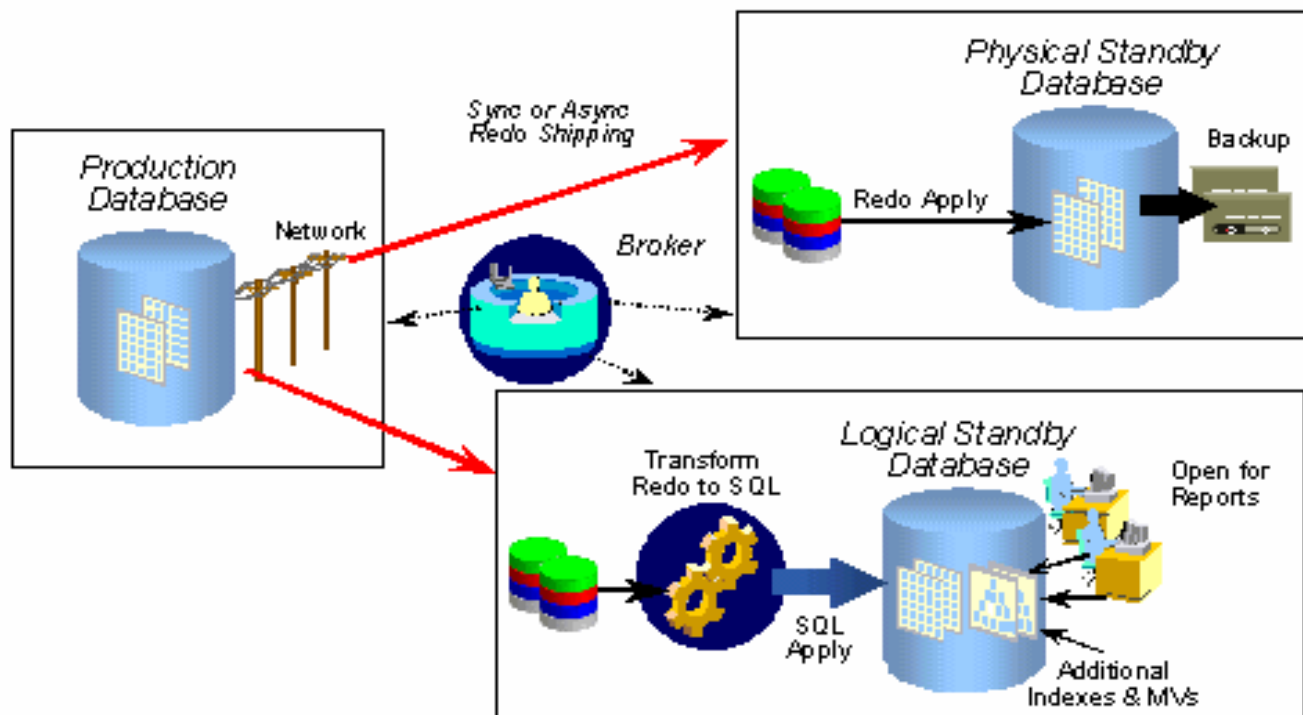


Oracle Data Guard

- Data Guard is Oracle's Disaster Recovery product which maintains and monitors one or more standby databases to protect enterprise data from failures, disasters, errors, and corruptions
- Standby databases, which can be located across large geographic regions away from the primary database, can be switched to the production role if a problem occurs with the primary
- Can use different Windows versions for primary and standby (2003 for primary, 2000 for standby)
- DG is free with Enterprise Edition of RDBMS

<http://www.oracle.com/technology/deploy/availability/htdocs/DataGuardOverview.html>



Data Guard Architecture



Maximum Availability Architecture (MAA)

- RAC + DataGuard
 - Eliminates physical location as SPOF

Data Guard





Patching Best Practices

- Apply CRS Patchset before RDBMS install
 - Simplifies process
 - Install patch to all nodes in the cluster
 - Apply patch to each node
 - Provides the ability for rolling updates
- For new installs on patched Oracle Homes, run catpatch after creating cloned databases with DBCA




Patching Best Practices

- For existing installs, make plans to test and apply future patchsets
- At minimum, CPU patches are released quarterly
- Fifth-digit patches provided on Windows platforms which bundle one-off patches
- For fixing new issues, applying latest patchset and fifth-digit patches available will help in eliminating known bugs



More Information

- Windows Server Technology Center
 - <http://otn.oracle.com/windows>
- Oracle Documentation on OTN: Windows tab at <http://www.oracle.com/pls/db102/homepage>
- <http://www.oratips.com> - debut edition in October contains article on Oracle/Windows by Scott Jesse
- “Oracle9i for Windows 2000 Tips & Techniques”
 - Authors: Scott Jesse, Matthew Hart, Michael P. Sale
- For more questions
 - karin.brandauer@oracle.com



The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remain at the sole discretion of Oracle.



ORACLE IS THE INFORMATION COMPANY

ORACLE®