



ESRI, 380 New York St., Redlands, CA 92373-8100 USA • TEL 909-793-2853 • FAX 909-793-5953

ESRI Enterprise Consulting Technical Brief

Oracle's Real Application Clusters and Transparent Application Failover with ESRI Products

Rev. 2
January 9, 2008

Copyright © 2007 ESRI
All rights reserved.
Printed in the United States of America.

The information contained in this document is the exclusive property of ESRI. This work is protected under United States copyright law and other international copyright treaties and conventions. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, except as expressly permitted in writing by ESRI. All requests should be sent to Attention: Contracts and Legal Services Manager, ESRI, 380 New York Street, Redlands, CA 92373-8100 USA.

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

U.S. GOVERNMENT RESTRICTED/LIMITED RIGHTS

Any software, documentation, and/or data delivered hereunder is subject to the terms of the License Agreement. In no event shall the U.S. Government acquire greater than RESTRICTED/LIMITED RIGHTS. At a minimum, use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in FAR §52.227-14 Alternates I, II, and III (JUN 1987); FAR §52.227-19 (JUN 1987) and/or FAR §12.211/12.212 (Commercial Technical Data/Computer Software); and DFARS §252.227-7015 (NOV 1995) (Technical Data) and/or DFARS §227.7202 (Computer Software), as applicable. Contractor/Manufacturer is ESRI, 380 New York Street, Redlands, CA 92373-8100 USA.

@esri.com, 3D Analyst, ACORN, ADF, AML, ArcAtlas, ArcCAD, ArcCatalog, ArcCOGO, ArcData, ArcDoc, ArcEdit, ArcEditor, ArcEurope, ArcExplorer, ArcExpress, ArcGIS, ArcGlobe, ArcGrid, ArcIMS, ARC/INFO, ArcInfo, ArcInfo Librarian, ArcInfo—Professional GIS, ArcInfo—The World's GIS, ArcLessons, ArcLocation, ArcLogistics, ArcMap, ArcNetwork, *ArcNews*, ArcObjects, ArcOpen, ArcPad, ArcPlot, ArcPress, ArcQuest, ArcReader, ArcScan, ArcScene, ArcSchool, ArcScripts, ArcSDE, ArcSdl, ArcSketch, ArcStorm, ArcSurvey, ArcTIN, ArcToolbox, ArcTools, ArcUSA, *ArcUser*, ArcView, ArcVoyager, *ArcWatch*, ArcWeb, ArcWorld, ArcXML, Atlas GIS, AtlasWare, Avenue, Business Analyst Online, BusinessMAP, Community, CommunityInfo, Data Automation Kit, Database Integrator, DBI Kit, EDN, ESRI, ESRI BIS, ESRI—Team GIS, ESRI—*The GIS Company*, ESRI—The GIS People, ESRI—The GIS Software Leader, FormEdit, GeoCollector, Geographic Design System, Geography Matters, Geography Network, GIS by ESRI, GIS Day, GIS for Everyone, GISData Server, JTX, MapBeans, MapCafé, MapData, MapObjects, Maplex, MapStudio, ModelBuilder, MOLE, NetEngine, PC ARC/INFO, PC ARCPLOT, PC ARCSHELL, PC DATA CONVERSION, PC STARTER KIT, PC TABLES, PC ARCEDIT, PC NETWORK, PC OVERLAY, PLTS, Rent-a-Tech, RouteMAP, SDE, Site-Reporter, SML, Sourcebook-America, Spatial Database Engine, StreetEditor, StreetMap, Tapestry, the ARC/INFO logo, the ArcAtlas logo, the ArcCAD logo, the ArcCAD WorkBench logo, the ArcCOGO logo, the ArcData logo, the ArcData Online logo, the ArcEdit logo, the ArcEurope logo, the ArcExplorer logo, the ArcExpress logo, the ArcGIS logo, the ArcGIS Explorer logo, the ArcGrid logo, the ArcIMS logo, the ArcInfo logo, the ArcLogistics Route logo, the ArcNetwork logo, the ArcPad logo, the ArcPlot logo, the ArcPress for ArcView logo, the ArcPress logo, the ArcScan logo, the ArcScene logo, the ArcSDE CAD Client logo, the ArcSDE logo, the ArcStorm logo, the ArcTIN logo, the ArcTools logo, the ArcUSA logo, the ArcView 3D Analyst logo, the ArcView Data Publisher logo, the ArcView GIS logo, the ArcView Image Analysis logo, the ArcView Internet Map Server logo, the ArcView logo, the ArcView Network Analyst logo, the ArcView Spatial Analyst logo, the ArcView StreetMap 2000 logo, the ArcView StreetMap logo, the ArcView Tracking Analyst logo, the ArcWorld logo, the Atlas GIS logo, the Avenue logo, the BusinessMAP logo, the Community logo, the Data Automation Kit logo, the Digital Chart of the World logo, the ESRI Data logo, the ESRI globe logo, the ESRI Press logo, the Geography Network logo, the MapCafé logo, the MapObjects Internet Map Server logo, the MapObjects logo, the MOLE logo, the NetEngine logo, the PC ARC/INFO logo, the Production Line Tool Set logo, the RouteMAP IMS logo, the RouteMAP logo, the SDE logo, The Geographic Advantage, The World's Leading Desktop GIS, *Water Writes*, www.esri.com, www.esribis.com, www.geographynetwork.com, www.gis.com, www.gisday.com, and Your Personal Geographic Information System are trademarks, registered trademarks, or service marks of ESRI in the United States, the European Community, or certain other jurisdictions.

Other companies and products mentioned herein may be trademarks or registered trademarks of their respective trademark owners.

Table of Contents

Background	4
Scope.....	4
Definitions	4
What Is Real Application Clusters?	4
What Is Transparent Application Failover?	4
Why Use RAC or TAF?	5
Implementation Considerations.....	5
Server-Side Considerations.....	5
Client-Side Concerns.....	6
Summary	6
Support	7
Suggested Resources	7

Background

As geographic information system (GIS) technology becomes more pervasive in business environments and in the enterprise IT realm, the concepts of high availability (HA) and scalability often arise as an area of discussion. One manner to address both these concerns is through the use of Oracle Real Application Clusters (RAC) for those customers whose IT standard or policy is utilizing Oracle as their RDBMS of choice.

Scope

The purpose of this paper is to introduce Oracle RAC, provide a high-level overview, and discuss how Oracle RAC may be used in conjunction with ESRI software to support an HA configuration. It does not however, go into detailed steps on how to set up ESRI's software with Oracle RAC.

This paper also provides a brief, high-level discussion of Oracle Transparent Application Failover (TAF), since Oracle TAF is often used in conjunction with Oracle RAC.

Definitions

What Is Real Application Clusters?

Real Application Clusters is a technology that links several individual computers so that they function as one. RAC enables each computer (or "node") that is a member of the cluster database to share access to the Oracle database. If one cluster node fails or is taken offline, the other nodes of the cluster continue operating, and the entire RAC database remains available for applications. This means that two or more relatively inexpensive computers appear to applications as if they were one much more powerful server. Oracle RAC utilizes what is called cache fusion. Cache fusion enables nodes on a cluster to synchronize their memory caches efficiently using a high-speed cluster interconnect so that disk I/O is minimized. Oracle RAC keeps track of all resources and resource states on each node and uses this information to restart these resources on opposite or surviving nodes.

What Is Transparent Application Failover?

Transparent Application Failover is a runtime failover for high-availability environments, such as RAC and RAC Guards, that refers to the failover and reestablishment of application-to-service connections. This reconnect happens automatically from within the Oracle Call Interface (OCI) library, which enables an application client to automatically reconnect to a database if the existing connection fails. Active transactions roll back, but the new database connection, which is achieved using a different node, is identical to the original. This is true regardless of how the connection fails. TAF reroutes application

clients to an available database node in the cluster when the connected node fails. Most application clients would not see error messages describing loss of service.

Why Use RAC or TAF?

Oracle's RAC addresses two key enterprise requirements for mission-critical database applications: high availability and scalability. It supports many different types of applications including data warehousing and online transaction processing (OLTP).

Additionally, more customers are leaning toward using low-cost commodity servers that run Linux and Windows to support their mission-critical database applications. Customers can save money when using RAC, which runs on low-cost servers instead of larger, expensive symmetric multiprocessing servers.

Moreover, customers may be able to take down a single node if and when any work needed to be performed on it, such as hardware maintenance or changes, and it could be done without affecting the overall availability of the application (assuming there are more than two nodes). Customers may also be able to add or remove nodes in real time, depending on their needs.

Oracle TAF, as a component of Oracle RAC, is really a feature of the Java Database Connectivity (JDBC) OCI driver. It enables the application to automatically reconnect to a database if the database instance to which the connection is made fails. When an instance to which a connection is established fails or is shut down, the connection on the client side becomes stale and will throw exceptions to the caller trying to use it. TAF enables the application to transparently reconnect to a preconfigured secondary instance, creating a fresh connection identical to the connection that was established on the original instance; that is, the connection properties are the same as that of the earlier connection. This is true regardless of how the connection was lost.

Implementation Considerations

Server-Side Considerations

For the most part, it has been a practice to avoid installing an ArcSDE service on any of the RAC cluster nodes. ESRI recommends using direct connect (DC) to establish connections to the Oracle database. There have been noted cases where ArcSDE services have been successfully implemented on clustered nodes; however, direct connect is still recommended because:

- The gsrvr process is dropped upon node failure.
- There is a service start-time delay upon failover to another node.

Some customers implementing RAC may also want to employ Oracle TAF, which is a component of RAC. The reason TAF may not work as expected is because ArcGIS uses cached cursors. Oracle TAF does not failover cached cursors. Therefore, any data edit sessions after the last save will be lost upon failover.

Other considerations include the time and correctly trained personnel to establish a production-level Oracle RAC configuration. Additionally, it is recommended by Oracle that a test environment be created and maintained as well. The test environment lends itself to further costs for hardware and maintenance purposes.

Client-Side Concerns

It is assumed that most applications in a clustered environment will not experience any lost connectivity when connecting to Oracle in an RAC environment. On the contrary, most clients, including ESRI's, are *not* cluster aware, making the assumption invalid. When a failover occurs, most clients (e.g., ArcIMS, ArcCatalog, and ArcGIS Server) will have to reconnect with the application. Since the user's transactions were most likely not completed, the user may simply have to reapply any unsaved work; that is, any work prior to the failover but subsequent to the last saved.

On the other hand, if it is a read-only session, Oracle TAF may work. ArcIMS as a client has the ability to reconnect once the service fails over. A known issue, (CQ00279163) was resolved as of release 9.2. Therefore, TAF should work if ArcIMS is the client; however, ESRI still recommends customers test this configuration.

Summary

ESRI has several customers who are successfully using ArcGIS products on Oracle RAC and Oracle TAF. These are observations that were made from previous experiences with Oracle's high-availability solution and ESRI software:

Advantages

1. High availability for instance failover
2. Active/Active server clusters with Oracle RAC for better scalability across nodes
3. The ability to "scale out" instead of "scale up," which provides a cost-effective hardware solution
 - a. Includes the use of inexpensive commodity servers
4. Reliability of RAC once established
5. More straightforward management of a large number of nodes

Disadvantages

1. There is additional hardware, licensing, and maintenance costs.
 - a. Includes the requirement for homogenous hardware, OS, and patch levels
 - b. Limits the environment to a single data center since it requires shared disk storage
2. Installation and administration are advanced.
3. Only saved transactions with TAF automatically failover as expected.

The main concern is how the client application will react within a clustered configuration. ArcIMS should be able to reconnect to the node after failover occurs. ArcCatalog and ArcMap clients are going to have to refresh their applications to make a reconnection. Other unknowns include ArcGIS Server as a client application and whether it reconnects as expected.

Support

Regarding support, the ESRI Technical Support and the Developer Support groups are not in a position to test or supply further information regarding Oracle high-availability solutions with the ArcGIS product suite. ESRI does not typically support third-party products directly. The information listed in this document covers only the most rudimentary questions.

Enterprise GIS system design is addressed in the *System Design Strategies* technical reference document at <http://www.esri.com/library/whitepapers/pdfs/sysdesig.pdf>. For answers to additional GIS capacity planning and solution questions, contact ESRI Systems Integration at sihelp@esri.com. For technical support, contact ESRI Technical Support at <http://support.esri.com>.

Suggested Resources

Beyer, Mark A., 2006. "Findings from Oracle OpenWorld 2006: Users Are Investing in Oracle's DBMS Infrastructure," Gartner, Inc., <http://www.oracle.com/database/docs/gartner-rac-asm.pdf>.

ESRI, "ArcSDE High-Availability Overview," http://www.esri.com/systemsint/kbase/docs/arcde_high-avail.pdf.

ESRI, "FAQ: Does ArcGIS support Oracle's Transparent Application Failover (TAF)?" <http://support.esri.com/index.cfm?fa=knowledgebase.techarticles.articleShow&d=27653>.

Oracle, MetaLink Note 97926.1, "Failover Issues and Limitations (Connect-time failover and TAF)."

Oracle, MetaLink Note 231996.1, "Quickstart Guide: Setup Transparent Application Failover (TAF) and Load Balancing."

Oracle, MetaLink Note 267909.1, "Real Application Cluster Knowledgebase Browser."

Yuhanna, Noel, 2005. "Oracle RAC Gains Momentum," Forrester Research, Inc., <http://www.oracle.com/database/docs/Forrester-On-RAC-0905.pdf>.