

Oracle Locator: Location-Enabling Every Oracle Database

*An Oracle White Paper
August 2005*

Oracle Locator: Location-Enabling Every Oracle Database

Introduction	3
Geometry, Layers	4
Spatial Indexing: R-trees.....	4
Spatial Operators	5
Fast Access With Two-Tiered Queries	5
Whole Earth Geometry Model for Geodetic Coordinate Support.....	6
Projections and Coordinate Systems – Support Based on EPSG Model (NEW in 10g Release 2)	6
Function-Based Index Support	7
Performance Enhancements.....	7
Enterprise Features Supporting Locator in Oracle Database 10g	8
Partitioning Support for Spatial Indexes	8
Parallel Spatial Index Creation	9
Parallel Spatial Queries (Introduced in 10g Release 1).....	9
Replication	9
Database Workspaces.....	9
Standard/Enterprise Edition Availability for Database Features.....	9
Open Standards	10
Support from Leading GIS and Location Services Vendors	10
Conclusion.....	11

Oracle Locator: Location-Enabling Every Oracle Database

INTRODUCTION

Every Oracle database includes powerful location data storage, query, and analysis capabilities. These native features of the database allow users to query location relationships from within the database, such as proximity of a first responder to an emergency call, sales by territory, and the location of service personnel relative to a maintenance site. Integrating location into business applications enables organizations to make better decisions, respond to customers more effectively, and reduce costs.

Oracle Locator, a feature of Oracle Database (Standard Edition, Standard Edition One, and Enterprise Edition), provides native data management, query, and analysis of location data, accessible through standard SQL. It provides a foundation for deploying location-enabled business applications, partner-based Geographic Information Systems (GIS) and location-based services. Developers can extend existing Oracle-based tools and applications, since with Locator they can easily incorporate location information directly in their applications and services. This is possible because location data is fully integrated in the Oracle server itself. Geographic and location data are manipulated using the same semantics applied to the CHAR, DATE or INTEGER types that are familiar to all users of SQL. Oracle Locator eliminates the cost of separate, proprietary systems, and is supported by all leading GIS vendors. It also brings the industry-leading security, performance, scalability, and manageability of Oracle Database 10g to critical location assets.

This paper describes Oracle Locator features¹, including new features introduced in the first and second releases of Oracle Database 10g.

¹ Oracle Spatial 10g is an option for Enterprise Edition that extends Locator with advanced spatial capabilities for geographic information systems applications and enterprise spatial information systems. Features include GeoRaster data management, network and topology data models, geocoding and routing engines, spatial analysis functions, and more. For more information, please refer to the *Oracle Spatial 10g Technical White Paper* and the *Oracle Spatial Option and Oracle Locator Data Sheet*.

For complete, detailed listings of Oracle Locator and Oracle Spatial features, please refer to Appendix B of the *Oracle Spatial User's Guide and Reference 10g Release 2 (10.2)*.

GEOMETRY, LAYERS

Oracle Locator supports three basic geometric forms that represent geographic and location data:

- Points: Points can represent locations such as buildings, fire hydrants, utility poles, oil rigs, boxcars, or roaming vehicles.
- Lines: Lines can represent features like roads, railroad lines, utility lines, or fault lines.
- Polygons and complex polygons with holes: Polygons can represent features like cities, districts, flood plains, or oil and gas fields. A polygon with a hole might geographically represent a parcel of land surrounding a patch of wetlands.

Internally, location data is modeled in layers, using a geometry column in a single table, sharing a common coordinate system. For example, the representation of a city might include separate layers for outlines of political districts or socioeconomic neighborhoods, every business and domestic location, and the maze of water, gas, sewer, and electrical lines. Because all these layers share a common database and notion of the Earth's geometry (coordinate, geoid and projection), they can be related through their respective (common) locations.

In addition to the geometric elements noted above, Oracle Locator supports the following geometry types:

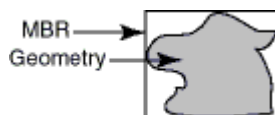
- Arc strings
- Compound polygons
- Circles
- Rectangles

SPATIAL INDEXING: R-TREES

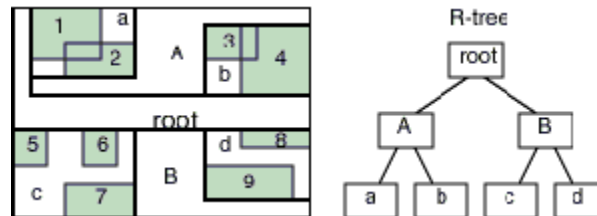
Oracle Locator applies spatial indexes - or **R-tree indexes** - to location data in an Oracle Database. R-tree indexes are simple to create, and require almost no tuning to achieve optimal performance. R-tree indexes can be created on two, three, or four dimensions of spatial data.

Typical queries specify a window of interest and retrieve all data intersecting or contained in the specified query window.

An R-tree index approximates each geometry with the smallest single rectangle that encloses the geometry (called the minimum bounding rectangle, or MBR).



For a layer of geometries, an R-tree index consists of a hierarchical index on the minimum bounding rectangles of all of the geometries in the layer. Because R-tree indexes are fast and work directly on geodetic data they are the preferred indexing mechanism for working with spatial data. Geodetic data is data consisting of angular coordinates (longitude and latitude) that are defined relative to a particular representation of the figure of Earth, or datum.



The spatial index uses the extensible indexing mechanism in Oracle Database 10g, providing maintenance operations of the index on insert, update, and delete. This results in increased ease of use.

SPATIAL OPERATORS

The interaction of various geometric features can be determined through the use of comparison operators, such as SDO_RELATE, SDO_CONTAINS, SDO_COVERS, SDO_ANYINTERACT (any interaction), and others. This permits answers to such requests as "list all the school zones crossed by this railroad line," or "find all pizza parlors within this area of interest." With Oracle Database 10g Release 1, new relationship operators were included as convenient alternatives to using SDO_RELATE with a mask value.

Locator also provides a function that computes distance between two geometry objects. This is useful for location-based services queries, such as "return the 10 hotels that are closest to the airport, and distance in miles to each." More advanced functions such as computing area or returning new geometries such as buffers, centroids, unions, intersections, or certain spatial aggregates require the Oracle Spatial option.

FAST ACCESS WITH TWO-TIERED QUERIES

Until now, database performance has largely been a factor of database size and index efficiency. But with Oracle Locator in Oracle Database 10g, performance is a function of the amount of data actually retrieved. Performance is optimized through the use of a spatial index and a two-tiered query model. This model significantly reduces load and query processing overhead and provides excellent scalability as the spatial data volume grows. The first tier, or primary filter, permits fast selection of a small number of candidate records to pass along to the secondary filter.

The primary filter uses approximations stored in the spatial index to reduce computational complexity.

The secondary filter applies exact computational geometry to the result set of the primary filter. These exact computations yield the final answer to a query. The secondary filter operations are more computationally intense, but they are only applied to the relatively small result set from the primary filter. Data can be returned to applications based on these approximations.

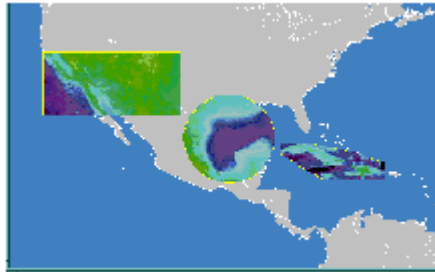


Image courtesy of TASC

Queries can be spatially constrained, as defined by an "area of interest" chosen by the user. Eliminating data outside the area of interest from consideration during queries ensures optimum performance levels.

Location queries, using standard SQL, can be made in a number of ways. For example, two-dimensional window extracts are possible as range searches, proximity searches, and polygon searches.

WHOLE EARTH GEOMETRY MODEL FOR GEODETIC COORDINATE SUPPORT

Oracle Locator provides a whole Earth geometry model that takes into account the curvature of the Earth's surface when performing calculations on geodetic data. The Oracle Locator distance function returns accurate distances for both projected and geodetic data (i.e., angular coordinates defined relative to a particular model of the shape of the Earth). Oracle supports over 30 of the most commonly used distance and area units which are useful for both geodetic and projected data, e.g. foot/square foot, meter/square meter, kilometer/square kilometer, and so on.

PROJECTIONS AND COORDINATE SYSTEMS – SUPPORT BASED ON EPSG MODEL (NEW IN 10G RELEASE 2)

Oracle Locator supports over 1000 commonly used mapping coordinate systems and also supports user-defined coordinate systems. It enables explicit map projection transformations of vector objects from one coordinate system to another. These transformations can be on a geometry-level basis or an entire layer (table) at a time.

In addition to the coordinate systems model provided with previous releases, in Oracle Database 10g Release 2, Oracle Locator provides coordinate systems support based on the European Petroleum Survey

Group (EPSG) data model and data set. EPSG support provides benefits of standardization, expanded support, and flexibility for oil and gas companies, georaster data vendors, and GIS users in general.

FUNCTION-BASED INDEX SUPPORT

A function-based index enables spatial queries and analysis on any relational data associated with a location attribute without creating and preloading a column of type SDO_GEOMETRY.

Users can create spatial indexes on location data stored in relational columns (for example in columns of longitude and latitude). Spatial operators can search function-based indexes as well as traditional spatial indexes. This spatial index will make it possible to invoke spatial operators on these relational columns without the need to create an SDO_GEOMETRY column.

This is useful for business geographic applications which have a schema for storing location data but cannot change their current schema to move the location data to a column of type SDO_GEOMETRY.

PERFORMANCE ENHANCEMENTS

Applications ranging from location-based services to GIS asset management and land management must frequently update and query location data — and demand high performance. With every new release, Oracle Locator has provided leaps in performance to address this requirement.

With the first release of Oracle Database 10g, Oracle Locator provided significant performance improvements over release 9.2:

- R-tree index inserts run 5-10 times faster
- R-tree index update time has been reduced by 40% or more - especially useful for enterprise geographic information systems and location-based services
- Spatial distance queries and “relate” queries run 20-40% faster
- Spatial joins run 2-6 times faster

With the second release of Oracle Database 10g, performance for R-tree index-based queries has improved.

Other Oracle Locator features can be used to boost performance further. For parallel queries on partitioned spatial indexes, performance scales with the number of CPUs used to execute the query. Building spatial R-tree indexes in parallel can dramatically reduce index creation time for very large non-point spatial datasets. Spatial aggregate functions speed retrieval of large sets of SDO_GEOMETRY objects. (Note: Locator includes the SDO_AGGR_MBR function; all other spatial aggregate functions are included only with the Oracle Spatial option.)

ENTERPRISE FEATURES SUPPORTING LOCATOR IN ORACLE DATABASE 10G

Oracle Database 10g provides powerful, reliable support for an organization's mission-critical applications. These enterprise features enrich Oracle's location capabilities via a flexible Internet deployment architecture, object capabilities, and robust data management utilities that ensure data integrity, data recovery, and data security. This level of support can only exist in the homogenous environment of an enterprise database solution, and cannot be effectively replicated in a hybrid solution that marries an external location-based solution with a traditional enterprise solution, no matter how tightly integrated the two components may appear.

Oracle Locator takes full advantage of expanded database size limits, high-performance VLDB maintenance utilities, replication, workspace manager (versioning), faster backup and recovery, and partitioning. Only users of Oracle's native spatial datatype can take full advantage of features such as partitioning, replication, parallel index builds and queries, and spatially-driven multi-level security. These features are not available or limited in functionality when using the LONG RAW or BLOB datatypes. The full range of Oracle utilities (e.g. SQL*Loader) are also available to ease migration and help upgrade applications that use the location-based services features. Some of these key enterprise features are described below.

Partitioning Support for Spatial Indexes

Oracle's database architecture includes partitioning, in which a single logical table and its indexes are broken up into one or more physical tables, each with its own index. Spatial indexes associated with partitioned tables can be partitioned; range partitioning is the partitioning scheme supported for spatial indexes.

Partitioning offers significant performance, scalability, and manageability benefits, including the following:

- Reduced response times for long-running queries; partitioning can reduce disk I/O operations.
- Reduced response times for concurrent queries; I/O operations run concurrently on each partition.
- Easier index maintenance, because of partition-level create and rebuild operations.
- Ability to rebuild indexes on partitions without affecting the queries on other partitions.
- Ability to change storage parameters for each local index independent of other partitions.

Partitions can also be split, merged, and exchanged.

Parallel Spatial Index Creation

Spatial indexes and index partitions can be created in parallel. R-tree index creation can be subdivided into smaller tasks that can be performed in parallel, making use of unused hardware (CPU) resources. For certain spatial data sets and index types and parameters, parallel index creation can substantially increase index build performance and provide a significant time savings. Large non-point datasets (commonly used in many standard GIS applications) can show dramatic performance improvements.

Parallel Spatial Queries (Introduced in 10g Release 1)

Spatial queries can now run in parallel on partitioned spatial indexes, improving the performance of "within distance", "nearest neighbor", and "relate" queries. Performance scales with the number of CPUs used to execute a query. This helps location service and land management applications, which need to execute high volumes of spatial queries quickly.

Replication

Oracle's Advanced Replication capabilities can be used for location data. For example, distributed systems that involve geographically dispersed yet logically replicated web sites, can take advantage of synchronized replication of spatial data objects across multiple databases.

Note: Advanced Replication multimaster configuration is offered with the Enterprise Edition database only. Refer to the *Oracle Database Advanced Replication* manual for more information about Advanced Replication features.

Database Workspaces

Oracle Workspace Manager, a feature of Oracle Database, provides a virtual environment (workspaces) that allows current, proposed and historical values for data to be managed in the same database. Workspaces can be shared and used to: isolate a collection of changes to production data until they are approved and merged into production; keep a long term history of changes to data; and create multiple data scenarios based on a common data set for "what if" analysis.

Standard/Enterprise Edition Availability for Database Features

Oracle Locator is available on Standard Edition, Standard Edition One, and Enterprise Edition. Some functionality requires core server features, which are not available or are limited on Standard Edition and Standard Edition One. Some of those features and their availability are listed below:

Feature	Standard/Enterprise Edition Availability
Parallel spatial index builds	Supported with Enterprise Edition only
Parallel spatial queries	Supported with Enterprise Edition only
Partitioned spatial indexes	Requires the Partitioning Option with Oracle Enterprise Edition. Not supported with Standard Edition.
Advanced Replication multimaster replication of SDO_GEOMETRY objects	Supported with Enterprise Edition only. (For more information, refer to Oracle Database Advanced Replication.)

For more on feature availability by edition, please refer to *Oracle Database Licensing Information* (part of the Oracle Database documentation).

OPEN STANDARDS

Oracle consistently works to help shape, drive, implement and support the latest open standards in the spatial and location services areas. Oracle is a Principal Member of the Open Geospatial Consortium (OGC) and participates actively on the Technical Committee. With Oracle Database 10g Release 1 (10.1.0.4), Oracle Locator complies with the OpenGIS Simple Features Specification for SQL, Revision 1.1, Types and Functions Alternative. Oracle is also committed to supporting the new OGC Geographic Markup Language (GML) as well as Open Location Service interfaces. The object-relational model used for geometry storage by Oracle Locator also conforms to the specifications associated with SQL92 representation of points, lines, and polygons.

SUPPORT FROM LEADING GIS AND LOCATION SERVICES VENDORS

Oracle Locator is directly integrated with the leading GIS and location services technology vendors. The breadth of partner support provides developers with their choice of best of breed tools to meet their requirements. With Oracle Database 10g and partner tools, developers can rapidly deploy scalable, secure enterprise GIS and location service solutions.

A list of partners is available at www.oracle.com/technology/products/spatial (click “Partners”, under “Quick Picks”).

CONCLUSION

Oracle Locator brings core location functionality to every Oracle Database, without requiring costly third-party extensions. Locator supports workgroup to enterprise deployments, and is ideal for many service based applications (such as location-based services). It allows organizations to use standard GIS partner tools to access the SDO_GEOMETRY type. It also allows any business application to include location capabilities to unlock the value of information they may already have. Locator supports spatial object type storage, SQL access, R-tree spatial indexing, spatial operations, and geodetic data storage and management, EPSG coordinate systems model, and more.

Oracle holds an 80-90% share of the geospatial database management market (IDC, *Oracle 10g: Spatial Capabilities for Enterprise Solutions*, Sonnen and Morris, Feb. 2005). Customers and partners rely on Oracle to deliver performance, scalability, security, and ease of use for their spatial applications. With Oracle Database 10g, Oracle Locator advances to new levels of performance and includes a rich feature set that makes it easy for any application developer to use Oracle to deploy GIS, location-enabled business applications, and wireless location-based services.



Oracle Locator: Location-Enabling Every Oracle Database

August 2005

Authors: Jean Ihm, Xavier Lopez

Contributing Authors: Jim Steiner, Jayant Sharma, Siva Ravada, Dan Abugov

Oracle Corporation

World Headquarters

500 Oracle Parkway

Redwood Shores, CA 94065

U.S.A.

Worldwide Inquiries:

Phone: +1.650.506.7000

Fax: +1.650.506.7200

oracle.com

Copyright © 2005, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle, JD Edwards, and PeopleSoft are registered trademarks of Oracle Corporation and/or its affiliates.

OGC, OpenGIS®, and CERTIFIED OGC COMPLIANT are trademarks or registered trademarks of Open Geospatial Consortium, Inc. in the United States and in other countries. Other names may be trademarks of their respective owners.