

# Oracle Spatial and Graph GeoRaster

ORACLE WHITE PAPER | NOVEMBER 2016





## Disclaimer

The following 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 remains at the sole discretion of Oracle.



ORACLE®



## Table of Contents

Disclaimer	1
Introduction	1
Architecture	3
Data Model	5
GeoRaster Object	6
Features of GeoRaster	8
Database Creation	8
Database Administration	9
Basic Data Manipulation	9
Raster Algebra and Analytics	10
Image Processing and Serving	11
GeoRaster Features in Oracle Database 12c Release 1 (12.1)	12
New Features in Oracle Database 12c Release 2 (12.2) on Oracle Cloud	13
Benefits of Managing Raster Data in Oracle Database	14
Conclusion	15



## Introduction

GeoRaster is a feature of Oracle Spatial and Graph that lets you store, index, query, process, analyze, and serve georeferenced raster image and gridded data and its associated metadata. It provides native data types and an object-relational schema to store and manage multidimensional array, grid layers and digital images that can be referenced to positions on the Earth's surface or in a local coordinate system.


What differentiates GeoRaster is the ability to perform raster analysis on extremely large images and data sets, provide in-place image processing and analysis with no development required, and provide parallelized image processing with simple invocation of PL/SQL procedures.

GeoRaster is used with data from any technology that captures or generates raster data and images, such as remote sensing, photogrammetry, and geospatial thematic mapping. It is used in a wide variety of application areas, including location-based services, geospatial image archiving and serving, environmental monitoring and assessment, geological engineering and exploration, natural resource management, climate modeling, ocean floor surveying and mapping, defense, emergency response, telecommunications, transportation, urban planning, and homeland security.

GeoRaster is designed to deliver enterprise-class data management capability to large image processing and GIS solutions and business applications. Developers can integrate this powerful data management technology with the leading image processing and raster/grid analysis tools and various business applications.

It meets the data management needs of broad application groups including:

- » Remote sensing, photogrammetry, GIS and geosciences applications – users manage and process their geospatial raster and gridded data assets using a scalable, secure, and robust RDBMS for defense, intelligence, agriculture, environment and natural resource management.
- » Business applications – leverage raster-based data in conjunction with other basic location data (address, etc.) to evaluate site locations and track fixed and/or continuous assets. These include Asset Management and Facilities Management, change analysis, site selection, suitability analysis and insurance risk assessment.
- » Image and Gridded Raster Data Repositories/Clearinghouses – support for clearinghouse servers that need to ingest, store, process, analyze, and disseminate very large volumes of geoinmagery and raster grids, through intranet or internet.



With Oracle Database 12c, Oracle Spatial and Graph GeoRaster provides substantial performance, data manageability, usability, and security features. GeoRaster supports virtual mosaics -- large collections of GeoRaster objects. Virtual mosaics are comprised of one or more GeoRaster tables (or views) and can be queried as a single GeoRaster object. GeoRaster raster algebra has analytical algorithms and faster, parallel raster operations. Image processing capabilities have been expanded and there is more server-side and parallel processing. The GeoRaster Java APIs have been expanded to support all PL/SQL APIs and also provide a new Virtual Mosaic Java API. GeoRaster now supports native JPEG 2000 compressions and has metadata to specify default alpha channel and pyramid level and supports a new resampling algorithm. GeoRaster includes a GDAL-based extract, transform and loading (ETL) tool that automates and enables concurrent batch loading and exporting of various image and raster files using GDAL.

With GeoRaster, it is possible to manage both georeferenced raster and spatial vector data in a single Oracle database. This means that the same storage, spatial referencing, indexing, and spatial operators can be used to store, query, and retrieve both raster and vector data. Moreover, all Oracle tools and utilities support the management of raster data and related datasets in the data server. For the customer, this means that the cost of building and fielding applications that use georeferenced raster image and gridded data can be reduced while the security, scalability and reliability of the application environment will improve dramatically.

This paper provides an introduction to the Oracle Spatial and Graph GeoRaster architecture, data model, object type, and major GeoRaster features. It also explains the benefits of managing raster data in Oracle Database.

## Architecture

The GeoRaster architecture provides the functionality needed to support the storage and use of image or grid-based raster data in Oracle Database. At a very high level of abstraction, the GeoRaster architecture includes five basic components:

- » GeoRaster Engine – This is the core, which provides the native GeoRaster object type and functionality including raster data and metadata indexing, update, query and manipulations.
- » PL/SQL API – Standard SQL access to the raster and grid-based data in GeoRaster databases.
- » Java API – Pure Java access to the GeoRaster objects in the databases. Users can also use JDBC, OCI, and OCCI to access to the GeoRaster objects without calling the GeoRaster PL/SQL or Java APIs.

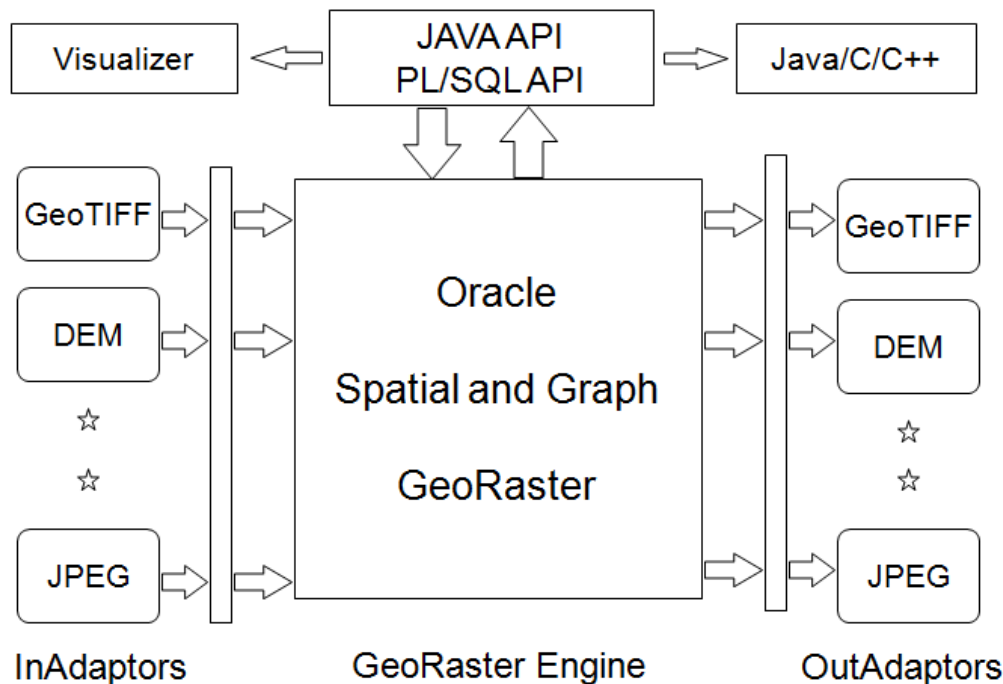


Figure 1. The Simplified Architecture of GeoRaster

- » Viewing Tools: Oracle Fusion Middleware MapViewer supports GeoRaster. A standalone viewer comes with the Oracle Spatial and Graph GeoRaster installation and can be used as a development or DBA tool. Finally, a variety of third party visualization and analysis tools support GeoRaster.
- » Input and Output [data] adaptors –GeoRaster includes an ETL tool based on the popular Geospatial Data Abstract Layer (GDAL) for concurrent batch loading and exporting of large numbers of raster files. This facilitates loading and unloading raster data between well-known image file formats and GeoRaster. The GDAL tool supports all raster formats supported by GDAL. GeoRaster also provides limited importing and exporting capability on six standard image file formats through both the server-side PL/SQL API and the client-side Java tool.

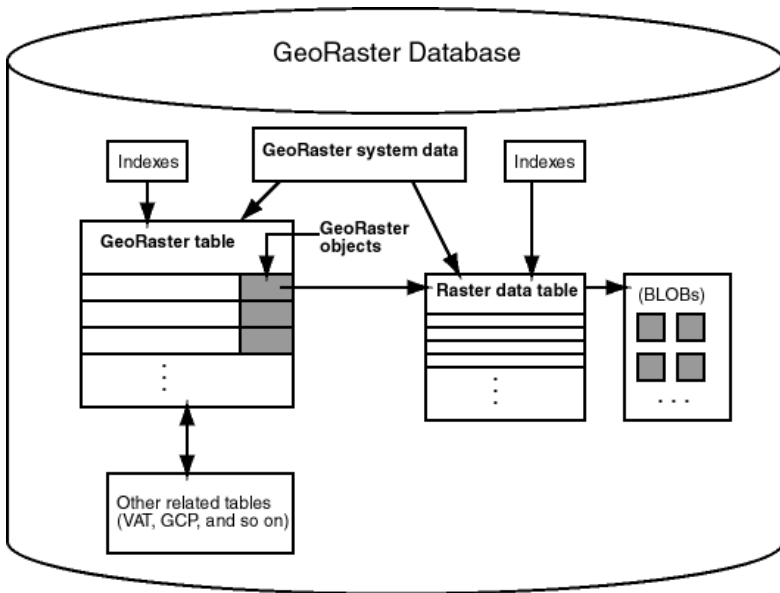


Figure 2. The Physical Storage Schema of GeoRaster

The core of Oracle Spatial and Graph GeoRaster is the physical schema designed to facilitate storing and managing raster or grid-based data inside the database. In the GeoRaster engine, the native data type called `SDO_GEORASTER` is defined and each image or raster grid is stored as a single object of this native type. A GeoRaster table is any user-defined table, which has at least one data column of type `SDO_GEORASTER`. `SDO_GEORASTER` objects include metadata and information about how to retrieve GeoRaster cell data that is stored in another user-defined table called a Raster Data Table (RDT), which is an object table of type `SDO_RASTER` or a relational table with the same columns as the attributes of type `SDO_RASTER`. The raster data table includes a BLOB column called `RASTERBLOCK`, which stores the real raster blocks. Other information associated with the GeoRaster objects can be stored in separate columns or tables, such as a Value Attribute Table (VAT). The relationship between a GeoRaster object and its raster data table is automatically managed by GeoRaster using a database dictionary. While users need to create and drop the raster data tables, they are internal to GeoRaster and the raster data in them are automatically updated and processed by the GeoRaster APIs.

A GeoRaster database basically consists of one or a list of GeoRaster tables, in which each image or raster grid is stored as one GeoRaster object in one row. It can contain an unlimited number of GeoRaster objects and each object can be terabytes in size. The GeoRaster tables can be in different database schemas and the GeoRaster objects can be accessed across schemas.

The specifics of the GeoRaster data model and how this architecture is implemented in Oracle Database are provided in the sections below.

## Data Model

There are two basic raster data types supported in GeoRaster: grid-based data and image data.

- » Grid-based, or gridded, data is a general term used for raster data such as digital terrain elevation, land use and land cover information, pollution concentration, geological information, and rainfall information. It is a rectangular grid of cells that are aligned to the X and Y-axes overlying an area. Each cell in the grid has the same size; this size is the resolution of the grid. Grid data stores the attribute values or attribute index values for each cell in the grid. If the index values are stored, different attribute values can be associated with the index values and are typically stored in a relational table, called Value Attribute Table or Raster Attribute Table.
- » Digital imagery is a specialized type of raster data. It is a two dimensional array (a matrix or grid) of regularly spaced picture elements (pixels). An image is created from optical or other spectral sensors, and is collected using a variety of technologies including satellite remote sensing and airborne photogrammetry. The size of the pixel is referred to as the resolution of the image. Digital images can be composed of many bands, referred to as multispectral or hyperspectral. Raster Attribute Tables can be associated with the image or its bands.

GeoRaster defines an integrated raster data model for these data types. Conceptually, it is component-based, logically layered, and multidimensional. The core data in a raster is a multidimensional array or matrix of raster cells. Each cell is one element of the matrix, and its value is called the cell value. The matrix has a number of dimensions, a cell depth, and a size for each dimension. The cell depth is the data size of the value of each cell and it applies to all cells. This core raster data set can be blocked for optimal storage, retrieval and processing. Pyramids (generalized, lower-resolution versions of the image – useful for fast retrieval in web applications) of the core raster data can be generated, stored and processed the same way.

The raster data is logically layered. The whole raster data is called the object layer or layer 0, and consists of one or more logical layers (or sublayers). For example, for multi-channel remote sensing imagery, the sublayers are used to model the channels or bands of the imagery. In GeoRaster, each sublayer is a two-dimensional matrix of cells that consists of the row dimension and the column dimension.

Besides the core matrix of cells, a raster data object has specific metadata associated with it. In GeoRaster, the metadata are organized into components that contain the following information:

- » Object information
- » Raster information
- » Spatial reference system information
- » Date and time (temporal reference system) information
- » Spectral (band reference system) information
- » Layer information for each layer



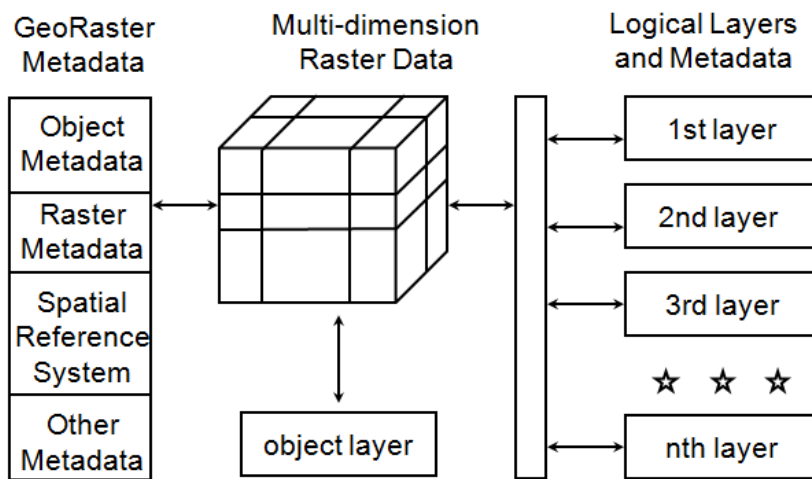


Figure 3. The Conceptual Data Model of GeoRaster Objects


The Object Information includes metadata such as user-defined ID, description and version information. The Raster Information includes metadata such as cell depth (e.g., 1BIT, 32BIT\_S, or 64BIT\_REAL), dimensionality, blocking sizes, interleaving types, compression and information about pyramids. The Spatial Reference System metadata contains information required for georeferencing, in which a generic polynomial model is defined and ground control points (GCPs) can be stored.

The Layer Information contains metadata pertaining to each logical layer in a GeoRaster object and it consists of many subcomponents. The major subcomponents of the Layer Information metadata include user-defined layer ID, scaling function, bin function, colormap, grayscale lookup table, statistics and histogram, NODATA values and value ranges, and bitmap masks. A Value Attribute Table can be used to maintain information about the values stored in each layer and the table name can be registered in the Layer Information metadata. Each layer, including object layer, can have one bitmap mask associated with it, which is registered in the Layer Information metadata. As special metadata, the bitmap masks are special raster grids with a cell depth of 1. The physical bitmap masks are stored together with the core raster data in its RDT.

Except Object Information and Raster Information, all other metadata components are optional. This enables GeoRaster to store, manage and manipulate generic multidimensional array data types as well.

## GeoRaster Object

Physically, the GeoRaster data model is embodied as two native RDBMS data types: SDO\_GEOASTER and SDO\_RASTER. The SDO\_GEOASTER data type encapsulates and stores all information about a raster data and is the external interface for users to create and manage a GeoRaster database. The SDO\_RASTER data type is internal and optional to the SDO\_GEOASTER data type and users don't need to directly manipulate it.



One raster data (an image or a grid) is stored in Oracle as an object of SDO\_GEOASTER data type, which is defined as:

```
CREATE TYPE sdo_georaster AS OBJECT (  
    rasterType NUMBER,  
    spatialExtent SDO_GEOMETRY,  
    rasterDataTable VARCHAR2(32),  
    rasterID NUMBER,  
    metadata XMLType);
```

The GeoRaster metadata is stored as the metadata attribute of the SDO\_GEOASTER type. It is an XML document using the Oracle XMLType data type. The metadata is stored according to the GeoRaster metadata XML schema defined by GeoRaster. The spatial extent (footprint) of a GeoRaster object is part of the metadata, but it is stored separately as an attribute of the GeoRaster object. This approach allows GeoRaster to take advantage of the spatial geometry type and related capabilities, such as using spatial R-tree indexing on GeoRaster objects and building huge global imagery databases. Another attribute of the SDO\_GEOASTER type is the rasterType, which contains dimensionality information and the data type that can be extended.


The actual raster cell data is internally blocked into small subsets for large-scale GeoRaster object storage and optimal retrieval and processing. All blocks are stored in a relational table or an object table of type SDO\_RASTER, which is defined as follows:

```
CREATE TYPE sdo_raster AS OBJECT (  
    rasterID NUMBER,  
    pyramidLevel NUMBER,  
    bandBlockNumber NUMBER,  
    rowBlockNumber NUMBER,  
    columnBlockNumber NUMBER,  
    blockMBR SDO_GEOMETRY,  
    rasterBlock BLOB);
```

This relational or object table is called raster data table, or simply RDT table. If the RDT is created as a relational table, it must have the same columns as the attributes of type SDO\_RASTER. Each block is stored in the RDT table as a binary large object (BLOB). Each row of the table stores only one block and the blocking information related to that block.

Both pyramid and bitmap masks of the GeoRaster object are stored in the same raster data table of the GeoRaster object, using the same blocking scheme. Pyramids can also be generated on the bitmap masks and stored in the same way as the pyramid of the GeoRaster object.

A blank GeoRaster object is a special type of GeoRaster object in which all cells have the same value. There is no need to store its cells in any SDO\_RASTER block; instead, the cell value is registered in the metadata in the blankCellValue element. GeoRaster also supports empty raster blocks to save storage space with large mosaic objects and to improve raster processing speed. It is used when there is no raster data available for a specific raster



block of a large GeoRaster object. Such GeoRaster data can be considered as a special sparse data type. There is still an entry in the raster data table for each empty raster block, but the length of the BLOB is zero, i.e., empty. Pyramids and bitmap masks can have empty blocks too.

The pair of the rasterDataTable and the rasterID attributes of the SDO\_GEOASTER object type uniquely identify a GeoRaster object in the whole database. Together, they provide the information internally required to store and retrieve the raster cell data in its raster data table. This design also allows many smaller GeoRaster objects share one RDT table while one huge GeoRaster object can flexibly have one devoted RDT table to improve performance and scalability. Users can have an unlimited number of GeoRaster objects stored in a GeoRaster table. Internally, GeoRaster uses a system dictionary (called the GeoRaster sysdata table) to maintain the relationship between GeoRaster objects and their related raster data tables. Even though stored in a separate RDT table, the raster cell data of a GeoRaster object is handled automatically by the GeoRaster functions, which are described in the next section.

## Features of GeoRaster

In addition to providing both a logical model and a physical model, GeoRaster provides the PL/SQL API and Java API with a rich set of foundation functions that facilitate data management for raster information in the Oracle Database. This section provides a general overview of the basic functional infrastructure available in GeoRaster.

GeoRaster operations can be grouped into the following categories:

- » Database Creation
- » Database Administration
- » Basic Data Manipulation
- » Raster Algebra and Analytics
- » Image Processing and Serving

### Database Creation

From a user perspective, a GeoRaster database is basically an Oracle database containing one or a list of GeoRaster tables, in which each image or raster grid is stored as a GeoRaster object in a row of a GeoRaster column. It can contain a virtually unlimited number of GeoRaster objects in one or more database schemas. One GeoRaster table can contain an unlimited number of GeoRaster objects and a single GeoRaster object can have only one pixel or be terabytes in size.

To build a GeoRaster database, users use the standard SQL statements or PL/SQL language and leverage the GeoRaster API and/or third party tools and solutions. The key steps include:

- » Create GeoRaster tables and RDT tables using the standard DDL.
- » Load raster data and metadata from files in various raster formats into the GeoRaster tables using third-party ETL tools, the GDAL-based ETL tool, the SDO\_GEOASTER.importFrom procedure, or the client-side Java loader. The SDO\_GEOASTER.importFrom procedure and the client-side Java loader are lightweight tools for conveniently load and export a limited number of image and raster files one at a time. The GeoRaster GDAL-based ETL tool is a powerful tool to load and export large numbers of raster and image files in batches and concurrently and it supports all raster formats supported by GDAL. The GDAL-based ETL and the Java loader are also integrated into the GeoRaster Viewer.
- » Generate and update the spatial extent on each GeoRaster object if necessary. The spatial extent is of SDO\_GEOMETRY type.

- » To validate the GeoRaster objects, call `SDO_GEOR.validateGeoRaster`. To validate the metadata, call `SDO_GEOR.schemaValidate`.
- » Build appropriate indexes on various columns of the GeoRaster tables e.g., a spatial R-tree index on the spatial extent attribute of the GeoRaster column and B-tree indexes on other columns. Users can also build function-based indexes on the metadata attribute of the GeoRaster objects.
- » The GeoRaster objects can be updated and deleted using standard SQL statements. They can be copied from one place to another using the `SDO_GEOR.copy` procedure.
- » To export the GeoRaster objects into files in various raster formats, use third-party ETL tools, the GDAL-based ETL tool, the `SDO_GEOR.exportTo` procedure, or the client-side Java exporter. Similarly as for loading, these tools are used to massively export large volume of raster data very quickly.

Once the database is created and the data are loaded into GeoRaster, users can manage it, tune it, perform spatial queries and various advanced processing as described in the next sections.

## Database Administration

The GeoRaster object type is a native Oracle data type. This approach allows users to manage the GeoRaster database using most standard RDBMS features, such as backup and recovery, partitioning, table security, replication, standby, and multitenency. GeoRaster supports Oracle Workspace Manager for raster block versioning and Oracle Label Security for row-level (raster block) security. In addition, GeoRaster provides the `SDO_GEOR_ADMIN` package, which includes more than 10 functions, to help manage and maintain GeoRaster databases. For example,

- » The `listGeoRasterObjects`, `listGeoRasterColumns`, `listGeoRasterTables`, `listRDT` subprograms in the `SDO_GEOR_ADMIN` package help check the status of existing GeoRaster objects and related objects in the current schema or the database.
- » The `SDO_GEOR_UTL.renameRDT` renames the RDT in the database to solve conflicts, which might happen during data migration.
- » The `SDO_GEOR_ADMIN.maintainSysdataEntries` automatically maintains the `SYSDATA` entries in the current schema or the database.
- » The `SDO_GEOR_ADMIN.upgradeGeoRaster` checks for and corrects errors after database upgrading.
- » The standalone GeoRaster Viewer is a DBA and development tool. It can connect to multiple databases or database schemas and display the GeoRaster objects. It allows zoom in and zoom out, scroll through the raster at any pyramid level and can query cell values and coordinates. It provides basic image processing operators and can view bitmap masks.

## Basic Data Manipulation

Besides leveraging the standard enterprise database features, GeoRaster provides over 200 raster data and metadata operations through the PL/SQL API to optimally manage and manipulate GeoRaster objects in support of various application requirements. Users can achieve many data management and manipulation goals using the following list of basic operations provided by GeoRaster:

- » Adjust the internal raster blocking size of any GeoRaster objects to improve scalability, optimize space usage and speed up raster processing and query. While each block must have the same size, the raster blocking sizes along different dimensions can be random numbers, not necessarily a power of 2.
- » Change the band interleaving types among BSQ, BIL and BIP to best fit into different applications.
- » Change the cell depth. Cell depths of 1-bit to 32-bit integers and 32-bit and 64-bit real numbers are supported.

- » Generate pyramids using different resampling approaches and remove pyramids. GeoRaster supports resampling methods of nearest neighbor, bilinear interpolation, biquadratic interpolation, cubic convolution, and average using four or six neighboring cells. The pyramiding operation supports parallelism to dramatically improve performance.
- » Compress or decompress GeoRaster objects natively in lossless DEFLATE, lossy JPEG compression, lossy or lossless JPEG 2000 compression types. MrSID compressions are supported through third-party plugins. Oracle Database Advanced LOB Compression can be used as well. Most GeoRaster functions that can be performed on uncompressed (decompressed) GeoRaster objects can be performed directly on compressed objects.
- » Georeference GeoRaster objects using PL/SQL functions or load georeferencing information from files. A generic polynomial georeferencing model is supported for georeferencing rectified and unrectified airborne photos and satellite images. It supports up to a power of 5 and 3-D model coordinates. The affine transformation, DLT, RPC, and other models are special types of this generic model. GCPs can be stored and used to georeference GeoRaster objects. The following commonly used geometric models are supported with GCP georeferencing: Affine, Quadratic Polynomial, Cubic Polynomial, DLT, Quadratic Rational, and RPC.
- » Add or remove bitmap masks for GeoRaster objects and individual bands or layers. These bitmap masks are stored inside the GeoRaster objects. Pyramids of masks can be created and stored inside the GeoRaster objects as well. Masks are compressed when the GeoRaster object is compressed.
- » Add or delete NODATA. Multiple NODATA values and multiple NODATA value ranges are supported for GeoRaster objects and their individual bands or layers.
- » Add or delete scaling function, bin function, colormap and grayscale information.
- » Spatially search GeoRaster objects using Oracle Spatial and Graph index and operators, such as area-of-interest queries, spatial join queries, and other topology-based spatial operations.
- » Crop images and perform subsetting to create new GeoRaster objects for persistent storage in the database or as an answer to a specific query for web distribution and display. Clipping the query or subset result along the polygon (irregular) boundary is supported.
- » Query cell coordinates and do coordinate transformation between GeoRaster cell space and model space. GeoRaster supports sub-cell or sub-pixel addressing (floating row and column numbers) in the GeoRaster cell spaces.
- » Query cell values or evaluate a direct location based on neighboring cell values by using a specified interpolation method, and returns the raster values for the specified bands or layers for that location.
- » Partially edit and update a window of raster data and its pyramids using another image or gridded data. Change the value of a single cell. The upper pyramid levels in the updated area can be partially regenerated to reflect the updated areas only.
- » Analyze statistics and generate histograms on the whole object or individual layers, both persistently or on-the-fly.
- » Query, delete and update most other items of the metadata through tens of functions or procedures. GeoRaster cell data and metadata update and query are crucial to successful use of GeoRaster in the Oracle Database. Some key subprograms include updating version information, querying the user-defined ID, the dimension sizes and blocking sizes, checking the coordinate system and georeferencing information, querying and updating time information, to name just a few.

## Raster Algebra and Analytics

GeoRaster provides a PL/SQL-based raster algebra engine enabling fast cell data searching, raster data analysis, and cartographic modeling. The GeoRaster raster algebra language is an extension to the PL/SQL language, which includes specific algebraic expressions and functions. The algebraic expressions support general arithmetic, logical, relational, and statistical operations. The raster algebra functions are mostly of local function type and include four major procedures to support arithmetic operation, conditional query, classify and cell value-based update. With this in-database raster algebra engine, user can do the following major operations.

- » Perform complex arithmetic operations.

- » Search cells based on logical and relational conditions.
- » Apply arithmetic operations to cells and then segment the resulting raster.
- » Edit or update cells of a raster based on conditions.
- » Generate statistics, including min, max, mean, median, mode, and standard deviation, on a layer or stack basis.
- » Perform complex logical operations and analysis.
- » Scale, offset and cast raster data.
- » Run cartographic models or other analytics.


These functions take one or many layers from one or many GeoRaster objects, apply the raster algebra expressions over those layers, do the specific algebraic computation or modeling, and output a new GeoRaster object or statistic results. The expressions can be defined in any way based the syntax of the expression language. The PL/SQL language is used to declare variables and constants, organize statements, and program the raster algebra processes and models.

In addition, all raster algebra functions are parallelized to further improve performance.

## Image Processing and Serving

GeoRaster provides some advanced image processing and serving capabilities. These include GCP georeferencing, reprojection, rectification, orthorectification, image scaling, stretching, masking, filtering, image segmentation, NDVI computation, Tasseled Cap Transformation, image appending, bands merging, large-scale advanced image mosaicking, and virtual mosaic. The operations described here are most commonly used to process and serving geospatial images, particularly raw satellite imagery and airborne photographs. However, those operations, just like the GeoRaster raster algebra, apply to all raster data types.

- » Enlarge or shrink images using nearest neighbor, bilinear interpolation, biquadratic interpolation, cubic convolution, or average resampling.
- » Reproject (transform) a GeoRaster object to a different projection or coordinate system, either on-the-fly for serving or persistently for storage.
- » Rectify 2-D georeferenced raw images.
- » Orthorectify 3-D georeferenced raw images with either an average height or a DEM.
- » Warp or affine transform images
- » Mask images physically using bitmap masks, in addition to store masks as pure metadata.
- » Enhance image colors by image stretching, normalization, equalization, histogram matching, and dodging.
- » Filtering images using standard or customized filters.
- » Segment images using raster algebra.
- » Compute NDVI (*Normalized Difference Vegetation Index*) and other VI's from images.
- » Apply TCT (Tasseled Cap Transformation) on images to analyze physical ground features of soil brightness, vegetation greenness, soil and canopy wetness, etc.
- » Merge or union layers from GeoRaster objects. Extract a single layer or a subset of layers from a GeoRaster object.
- » Append smaller images to a large image without spatial boundary limitation. The smaller images can be rectified, un-rectified, or have different resolutions. They can be located on any side of the target image.
- » Mosaick large raster datasets inside the database. The GeoRaster mosaic function allows gaps, overlaps, and missing source GeoRaster objects. It supports both rectified and un-rectified images. It supports internal reprojection/rectification and common point rules. It supports color balancing including histogram matching by



automatically identifying and using overlapping areas. It is also implemented with parallelism. This mosaicking process results in a single GeoRaster object, which is also called a physical mosaic.

- » Support Virtual Mosaic, which is any large collection of georeferenced GeoRaster objects, rectified or un-rectified, from one or more GeoRaster tables or views that is treated as if it is a single GeoRaster object. A virtual mosaic is defined as one or a list of GeoRaster tables, a database view with a GeoRaster column, or any SQL query statement that results in a collection of GeoRaster objects. Pyramids of virtual mosaic are supported. Users issue a single call to query the virtual mosaic based on area-of-interest (i.e., subsetting or cropping) and request the cropped images to be in different coordinate system with different resolutions. On-the-fly transformation with resampling and mosaicking with common point rules and color balancing are done internally and automatically.
- » Image and raster data serving to clients or applications are supported through many functions in the PL/SQL API and Java API described in this paper. These include direct exporting to files, searching and then subsetting or cropping the rasters, and applying reprojection and rectification on-the-fly. Particularly, when the user doesn't want to create large physical mosaics, the virtual mosaic is best used in serving image database out to various applications. It doesn't require the source images to be preprocessed or mosaicked beforehand. Instead, all images are stored as is and the whole image data set can be served based on small areas-of-interest using single calls to the server.
- » Image and raster data can also be served over the web using standard OGC web services, i.e., WMS and WCS.

In summary, the PL/SQL API and Java API are provided for database creation, database administration, data manipulation, raster analysis, image processing, and raster data serving and delivery. They also aid in the development of and integration with applications on top of GeoRaster. Users can use Java, C or C++ to leverage the PL/SQL API and the Java API or directly access the binary data of the open GeoRaster data model. In addition, WMS and WCS are provided to serve image and raster data over the web. Many of the functions discussed above are extended, augmented or leveraged by partner technologies delivered as load/transform/export tools, comprehensive raster analysis and image processing client tools, or in the form of visualization engines built on top of the GeoRaster model.

## GeoRaster Features in Oracle Database 12c Release 1 (12.1)

This section outlines some of the features of GeoRaster in Oracle Database 12c Release 1, most of which are described in more detail in previous sections.

### Raster Algebra and Analytics

- » Provide a Raster Algebra Language, which is an extension to the PL/SQL language. It includes an expression language and 4 functions of Map Algebra local function type.
- » Support cell value based conditional queries using raster algebra.
- » Support cell value based updates (raster editing) using raster algebra.
- » Support any arithmetic operations using raster algebra.
- » Support classify for raster segmentation using raster algebra.
- » Seven statistics generation functions to support interactive analysis on-the-fly.
- » The existing generateStatistics is enhanced to support on-the-fly analysis.

### Image Processing and Virtual Mosaic

- » Support rectification of georeferenced raw images.
- » Support orthorectification of georeferenced raw images with height or DEM.
- » Support large-scale image appending.
- » Support large-scale and advanced mosaicking of georeferenced raw images.
- » Support Virtual Mosaic.

- » Both advanced mosaicking and virtual mosaic supports internal reprojection/rectification, 8 different common point rules, and simple color balancing.
- » Support image masking, image stretching, and image segmentation.
- » Support NDVI Computation and TCT.

#### **Parallel Processing and Enhancements to the GeoRaster Core**

- » Parallel processing is one of the major features in this release. Parallel processing capabilities are added into mosaicking, pyramiding and all raster algebra functions.
- » Operation Progress Reporting is added into many functions.
- » Significant performance improvement of pyramiding and mosaic.
- » Supports relational RDT tables as an additional raster data storage.
- » Users can specify default alpha channel and default pyramid level in the metadata.
- » Biquadratic interpolation method to interpolate cell value for any point in the grid. The biquadratic interpolation method is also added as a resampling type in image operations and pyramiding.
- » Added "OTHER" resampling type to indicate an unknown or external pyramid resampling type imported from files.

#### **Other Server-Side Functions**

- » Supports resolution generation for pyramids.
- » A function to compute raster range, i.e., the min and max cell values.
- » A template function to get lob locators.
- » Create empty raster blocks from existing raster blocks.
- » A DML trigger recreation utility function.

#### **Enhanced Client-Side Tools and Java API**

- » A GDAL-based ETL tool for concurrent batch loading and exporting of large numbers of raster files.
- » The GeoRaster Viewer is enhanced to display a virtual mosaic defined as one or a list of GeoRaster tables or views.
- » The Java API is enhanced to support features such as ground control point (GCP) storage and manipulation, GCP georeferencing, reprojection, and grid interpolations.

### **New Features in Oracle Database 12c Release 2 (12.2) on Oracle Cloud**

This section outlines some of the new features of GeoRaster in Oracle Database 12c Release 2.

- » Native JPEG 2000 compression is supported.
- » Mosaicking supports statistic matching and histogram matching color balancing approaches by automatically identifying and using overlapping areas.
- » A set of new Image processing operators including image stretching, equalization, normalization, histogram matching, dodging, and filtering.
- » Image warping and affine transformation are supported.
- » Over 25 new operators are added to the raster algebra language, including conditional (IF-THEN-ELSE) and statistical operators.
- » Stack based statistical and logical raster algebra functions are supported.
- » A set of new memory control functions to allow users to control how much memory to use during GeoRaster operations. Memory size can be set to improve processing performance.
- » New functions for and significant performance improvement on single cell value queries.
- » New functions to compute area weighted statistical mean value for the cells and sub-cells within a specific window, and to calculate the three-dimensional surface area represented by DEM.



- » New color ramping functions to generate a color gradient or grayscale gradient for grids and images.
- » All server-side new functions and procedures are supported in the GeoRaster Java API.
- » The GeoRaster ETL and Viewer are integrated into one tool for ease of deployment and use.
- » OGC Web Coverage Services (WCS) 2.0.1 is supported.


In summary, GeoRaster in the Oracle Database 12c release enhances raster data queries, processing, and serving with content based raster analytical queries, context based image aggregations, other new functions, and delivers significant performance improvements. It provides image processing and raster algebra functionalities, enabling massive in-database processing and analytics. Its high-performance parallel processing capabilities and many algorithms improve GeoRaster performance dramatically. Further details can be found in the Oracle Spatial and Graph GeoRaster Developer's Guide.

## Benefits of Managing Raster Data in Oracle Database

This section describes some of the major benefits of managing geospatial imagery and raster data using GeoRaster.

The GeoRaster data model is an open and integrated raster data model. It supports various image and raster data types and removes the need of dealing with many data models and significantly simplifies the data management tasks.

- » The GeoRaster data type is designed for raster data storage inside the database and is independent from any file formats. Raster files in different file formats are loaded into GeoRaster for efficient storage and management. All metadata and cell data of GeoRaster objects are open to users and can be accessed at bit level as long as the users are granted the access privilege. This enables shared third-party application developments as well as avoids vendor lock-ins.
- » It breaks the size barriers. A virtually unlimited number of GeoRaster objects can be stored and a single GeoRaster object can be multi-terabytes in size. The internal tuning tools and compressions can be used to optimize the raster blocking and other storage options to best fit into the needs and performance requirements of different applications.
- » The native data type approach provides great flexibility to store raster data in regular relational tables. Images or rasters in different projections and locations can be stored in the same table and users can easily build global databases with whole-earth spatial index and other indices enabling fast queries and manipulations of rasters anywhere on the Globe.
- » SQL access provides better management of raster data. The standard SQL or PL/SQL API is flexible, power and easy-to-use. It helps speed up enterprise integrations and broadens geospatial application areas.
- » The database-centric approach and balanced memory management allows thousands of concurrent queries and users. And the concurrency can scale-up drastically by leveraging Oracle Enterprise GRID computing technologies or using computer clusters.
- » High performance computing is made possible for raster data management and processing. GeoRaster features in-database processing, optimized multi-tasking, and parallel processing. The in-database data manipulations, raster algebra and image processing capabilities allow data to be processed where it is stored. Coupled with multi-tasking and parallel processing, this provides great performance and true security. Traditional image processing and GIS systems can leverage it to pre-process and filter data inside the database thus improve performance too.
- » Data security, replication, partitioning, bulk load utilities are readily available.
- » Raster, vector, XML and various types of attribute data can be stored on a single server. By effectively managing raster, vector and attribute data in a single data management environment – with common storage, indexing, spatial referencing, query optimization, security, and user management – Oracle reduces the processing overhead and eliminates the complexity of coordinating and synchronizing disparate sets of spatial data.

- 
- » It provides the best platform for consolidation of disparate data management environments, such as GIS, remote sensing, business data. Complete and customized data management solutions can be built for both business and traditional geospatial applications.
  - » Oracle Fusion Middleware MapViewer is tightly integrated with GeoRaster. Internet deployment enables a large number of concurrent users to access an application at virtually no additional cost with 24 x 365 uptime.
  - » It is supported by and integrated with leading 3rd party and open source image processing, GIS, web services and visualization tools.
  - » Reduced training, software, support and application integration costs resulting from consolidating raster, vector and attribute storage.
  - » Risk reduction – imagery and raster information are integrated into Oracle Database resulting in scalability, reliability, and fast performance.

## Conclusion

The availability of Oracle Spatial and Graph GeoRaster in Oracle Database 12c creates significant capacity for managing, processing, analyzing, and serving large volumes of raster data. Oracle is the only provider of commercial database management software that can store geospatial imagery and grid-based raster data in the database as a named native data type. It enables full integration of raster datasets with other enterprise datasets and supports better business applications. The open and integrated raster data model supports various image and raster data types and simplifies data management tasks. It allows flexible blocking of raster data so that large-scale rasters can be stored and easily managed. Spatial and other indices provide fast searching and retrieval of raster metadata and cell data. The GeoRaster PL/SQL and Java API help users optimize the internal storage of any existing GeoRaster objects, manage the database, search rasters based on area-of-interest, create large mosaics, process massive images, and conduct cartographic modeling with parallelism and on a global basis. Selective raster query tools and other operations can be easily built upon GeoRaster to quickly serve imagery and rasters to middleware and applications. It reduces the cost of building and fielding of applications. It is now possible for users of file-based image processing systems and raster data applications to benefit from the scalability, security, reliability and performance of Oracle Database and take advantage of enterprise grid computing technologies to support mission critical applications and integration.







**Oracle Corporation, World Headquarters**

500 Oracle Parkway  
Redwood Shores, CA 94065, USA

**Worldwide Inquiries**

Phone: +1.650.506.7000  
Fax: +1.650.506.7200

CONNECT WITH US

-  [blogs.oracle.com/oracle](http://blogs.oracle.com/oracle)
-  [facebook.com/oracle](http://facebook.com/oracle)
-  [twitter.com/oracle](http://twitter.com/oracle)
-  [oracle.com](http://oracle.com)

**Integrated Cloud Applications & Platform Services**

Copyright © 2016, Oracle and/or its affiliates. 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 and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0116

Oracle Spatial and Graph GeoRaster  
November 2016  
Qingyun (Jeffrey) Xie



Oracle is committed to developing practices and products that help protect the environment