

# Raster Data - Poor Second Cousin to Vector No More

By: [Joe Francica](#)

Directions Magazine (Jun 18, 2004)

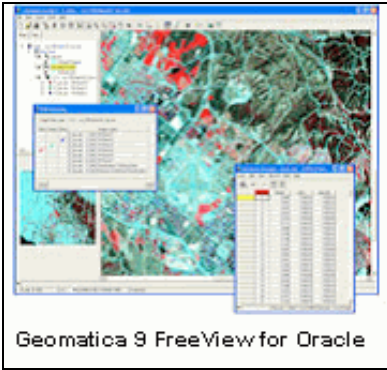
Over the course of the last ten to fifteen years, raster data has played the role of a poor second cousin to its vector brethren. As a technology community, we have tended to worship at the 'vector altar' because the data structure was more easily stored, retrieved and managed, and its graphical representations were more precise than the pixelated raster. And yet, as a structure for modeling spatial information, raster offers so much more in terms of quantitative analysis. But, pray all, raster is a storage hog, performance can be a big problem, and flat file management is devilishly difficult to control.

For those of you who remember the early days of image processing, whether it be for Landsat imagery or digital elevation models, flat file storage was the only choice you had. An algorithm would be employed to manipulate an image, and the results were captured as a separate image layer. This layer could then be employed as an input to yet another analysis but intrinsically, it was likely that some of the metadata was lost, except perhaps in handwritten notes. If you were doing multispectral image classification, the user generated many, many image layers. [How many of you used the old IDIMS system? Interactive Digital Image Manipulation System from ESL! Or the GE 100?]. The result... a storage nightmare. Back in those days, you didn't have huge hard disks. Image data ate 9-track tapes like whales eat krill.

Lately, things have changed for brother raster. The release of Oracle 10g is changing the way that raster data are managed. Previously, the Oracle Spatial module handled only vector data types. With Oracle 10g's GeoRaster feature for Oracle Spatial, a new data type, SDO\_GEORASTER, is defined, and raster data is stored as a new data table (an object table of type SDO\_RASTER). This is allowing raster manipulation and analysis software solution products to concentrate on what they know best - analysis - and permit Oracle to manage massive arrays of grid data. "10g is a huge deal," says Jeff Hoedeman, Director of Marketing for PCI Geomatics, a company that recently announced its support of Oracle GeoRaster with a product called GeoRaster ETL for Oracle.



At the core of Oracle's GeoRaster feature is the GeoRaster Engine that includes functionality for metadata management, methods, and indexing. It includes an SQL Application Programming Interface (API) and other programmatic interfaces using the Oracle Call Interface (OCI), Java, C and C++. According to Oracle, "The GeoRaster metadata contains XML representations of metadata that describes the raster data. This metadata includes object metadata such as descriptions, and version information. GeoRaster also includes raster metadata for cell depth, dimensionality, blocking, interleaving, and other information. Additionally, Spatial Referencing System metadata containing information for the affine transformation required for georeferencing can also be stored. "This and other details about the GeoRaster architecture can be found in a white paper entitled, "[Oracle Database 10g: Managing Geographic Raster Data using GeoRaster](#)."



Now that one huge headache for data management as been removed for both users and companies who develop raster processing software, the focus turns to manipulation and analysis. As mentioned above, PCI Geomatics recently announced their extract, transform, and load (ETL) product, GeoRaster ETL for Oracle. The objective of this product is to provide a "gateway" for loading raster data into Oracle 10g. In addition, GeoRaster ETL supports image metadata to map to Oracle 10g. GeoRaster ETL also takes across associated ancillary information such as color look up tables. For more information, [click here for the GeoRaster ETL website](#).

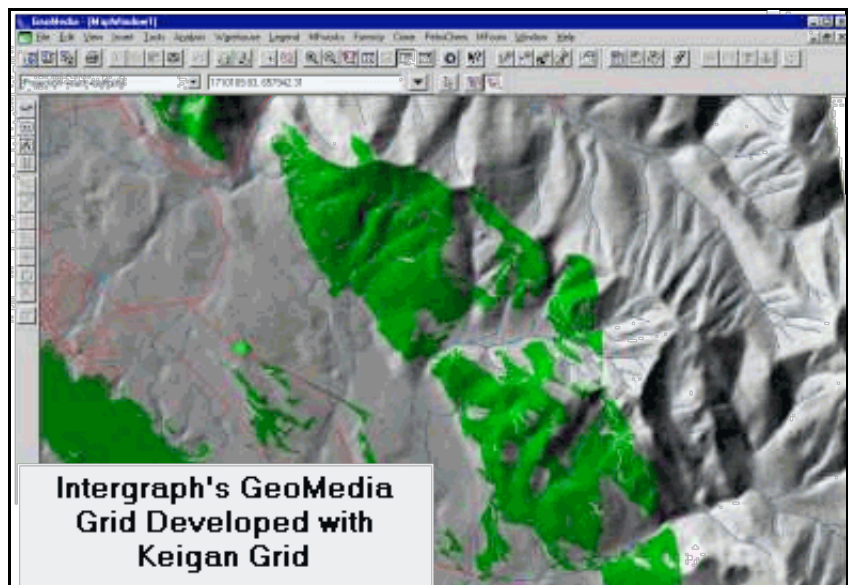
In terms of grid analysis, few products match the breadth of functionality like Keigan Systems' [Keigan Grid](#). Keigan adapted Dana Tomlin's Map Algebra suite of algorithms for grid data to create the original "Map II" product that later underwent several name changes before its latest incarnation as Keigan Grid. "The number of people who see the power associated with raster is on the increase," says Robert Gammon, Vice President of Keigan's Client Services.

Keigan Grid is a set of COM compliant objects for specifically integrating grid processing into the vector world. Keigan has created a number of controls to view, create legends, change colors, and translate data from grid to vector, and vice versa. The key to Keigan's product is the plethora of raster map operations for grid analysis. And because Keigan Grid conforms to the COM architecture, the user is allowed to select the pieces of that technology to build single executions for a customized solution. Keigan can now share its technology with other development organizations to embed grid technology in vector systems. This is what the company has done with Intergraph to create [GeoMedia Grid version 5.2](#).

Keigan has also announced support for Oracle 10g. Keigan will use a combination of the Oracle Call Interface (OCI) and Oracle C++ Call Interface (OCCI) to gain access into the database. This gives them the ability to use the database structure, which Oracle has defined, for storing raster data directly into the database. GeoMedia Grid for Oracle is slated for release this November.

In addition, other software providers are working on support for Oracle 10g. ER Mapper is working on full spatial imagery integration with databases, such as Oracle 10g and IBM's DB2, that will include image preparation, use, serving and integration capabilities.

Some performance issues remain with Oracle 10g's handling of very large databases. There is no compression of the raster data within the database at this time. Keigan, for example, has as compression flag associated with their data structure for GeoMedia Grid so that they can store the data in Oracle Spatial as either Keigan compressed version or a native



version. If the flag is indicated for the native version, the data will be stored as Oracle needs them to be stored in case other software tools that

have built an API to Oracle GeoRaster need to access to them.

So, although some issues remain, the ability to read, manage, and manipulate raster data has taken a quantum leap forward. With the volume of satellite data that is being captured, the higher pixel resolution, and the demand for accurate, near real-time data, the need to process and store raster information will only grow. The raster/vector divide is closing and integrating these data types is certainly a "prayer answered."

---



Reprinted with permission, Directions Media, Copyright 2005