

Managing Rasters in a Personal Geodatabase

By Chris Wayne, Mazama Mapping

As many users know, the most efficient and effective way to store large amounts of raster data is in a geodatabase. With ArcGIS 8.x, that geodatabase had to be a relational DBMS managed through ArcSDE. However, ArcGIS 9 provides powerful new tools for managing raster data within the personal geodatabase. This article describes how to use these new features.

Overview of Rasters in the Personal Geodatabase

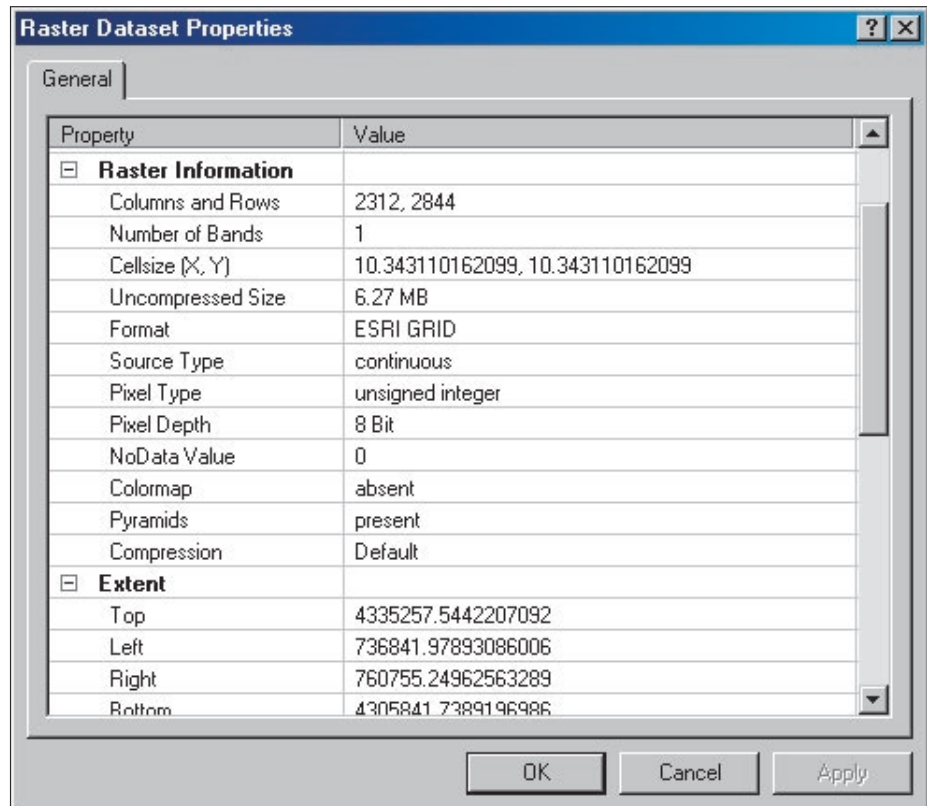
Within a personal geodatabase, rasters may be stored as raster datasets or catalogs. Raster datasets represent a continuous single raster. Catalogs are collections of individual rasters. Each storage method has strengths and limitations, but there are several general requirements that apply to both storage methods.

Raster catalogs and datasets can be created as empty containers before they are populated and subsequently loaded, or these geodatabase raster storage types can be created on the fly. Creating and loading raster datasets and catalogs are performed in ArcToolbox using tools in the Data Management Tools toolbox. These tools are exposed via the ArcCatalog user interface.

When creating a raster catalog, the XY domain (the combined spatial extent of all the input rasters) must be set; however, the XY domain is not applicable to raster datasets. As with all datasets within ArcGIS, it is highly recommended that raster data brought into either a catalog or a raster dataset have a spatial reference previously defined, either a Geographic Coordinate System (GCS) or a Projected Coordinate System (PCS). That reference does not necessarily need to be the same as the catalog or dataset.

Raster datasets have one spatial reference, and during the mosaic process, the pixels in different coordinate systems will be projected on the fly into the correct location for that particular raster dataset in the personal geodatabase. In the case of a raster catalog, individual raster items can have their own spatial reference that is different than the spatial reference of the geometry or raster columns, and these rasters will simply be projected on the fly for display or analysis.

Unlike rasters maintained by ArcSDE, which are stored within the underlying DBMS, rasters managed using a personal geodatabase are not actually stored in the personal geodatabase.



The Raster Dataset Properties dialog box in ArcCatalog features a tree control that makes it easy to quickly browse through properties for a raster and access options such as building statistics.

Whether as a catalog or raster dataset, rasters managed by a personal geodatabase are stored in a referenced location. Because only the reference is stored in the personal geodatabase, the actual raster size is not constrained by the two gigabyte storage limitation of the personal geodatabase. However, it is highly recommended that raster data referenced by a personal geodatabase be manipulated (i.e., copied, updated) with ArcGIS. Using Windows Explorer to manipulate this raster data will generate problems similar to those caused by managing coverages or GRIDs outside of ArcGIS.


Raster Datasets or Raster Catalogs?

Deciding whether to store rasters as a raster dataset or a raster catalog is actually the second decision to make. The first decision should be to determine if the raster data is suitable for use in a GIS and/or storage in a personal geodatabase. Taking the time to georeference raster data to real-world coordinates using the Georeferencing Toolbar or, perhaps better, obtaining data from a

vendor that provides the already georeferenced imagery is highly recommended. Most imagery comes with internal georeferencing or an associated world file.

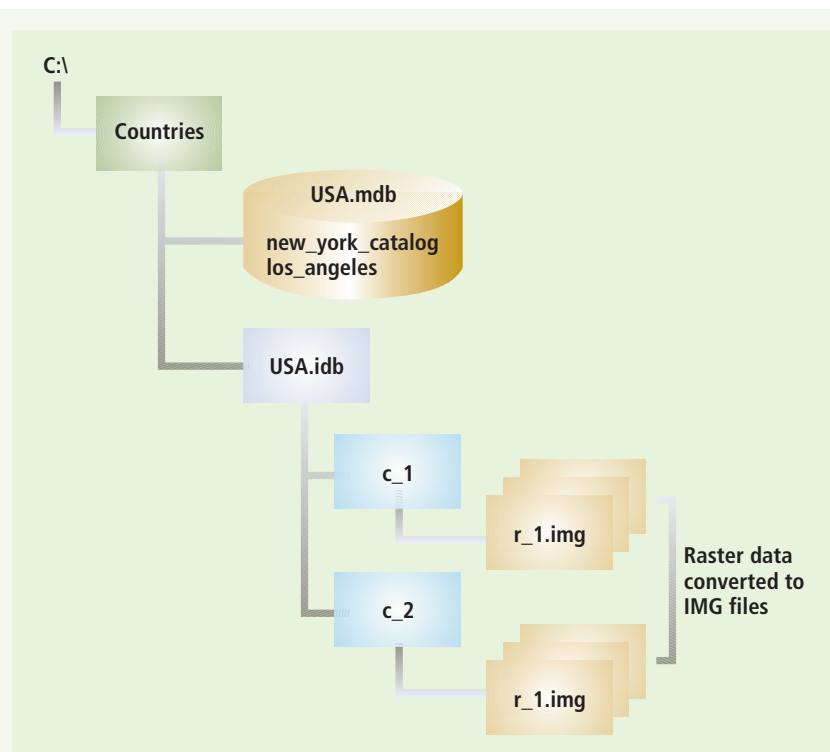
A GIS assumes all data is georeferenced. If rasters are not georeferenced, as is the case with ground or oblique photographs or scanned documents, using them in a GIS will cause problems. The images will appear spatially on top of each other in image space. The best way to work with "pictures," such as scanned images or photographs, is to associate them as raster attributes to real-world points, lines, or polygon feature classes. This is the third type of raster storage type within a personal geodatabase.

Having opted to use the personal geodatabase, the next step is to decide whether rasters will be stored as a single mosaicked raster dataset or individually within a raster catalog. Generally speaking, if these rasters will be used for analysis as a single entity, create a raster dataset. ArcGIS Spatial Analyst tools treat raster datasets like any other supported raster. (Note: All raster datasets within the personal



Nonmanaged Geodatabase Raster Data

- Only raster catalogs and rasters as attributes of feature classes can be nonmanaged.
- Raster data is referenced to file-based rasters of any supported raster format through the RASTER field.
- The path name is used to locate the raster file.
- Data is not duplicated/copied/converted.
- Rasters are managed by the user, not the geodatabase.
- Nonmanaged catalogs load more quickly, because there is no copying of raster data.



Managed Geodatabase Raster Data

- All raster datasets must be managed.
- Raster catalogs can be managed.
- Rasters as attributes of feature classes can be managed.
- Rasters are copied and stored in a folder at the same directory level as the personal geodatabase.
- Raster data is converted to IMG format files.
- Raster data is referenced to file-based IMGs through the RASTER field.
- Data is managed by the geodatabase, not the user.
- For rasters larger than 2 GB, an IMG file is created for the header information and IGE files store the actual imagery. In special cases in which the pyramid file (RRD) also grows larger than 2 GB, an RDE file is created to hold this information.

geodatabase are of the “managed” type.)

There are variations in raster datasets. A dataset may be a single raster that is imported from another source, such as an ESRI Grid, ERDAS IMAGINE IMG, or a TIFF/GeoTIFF image. Alternatively, multiple rasters can be added to a dataset to make a single mosaic from many different rasters. This process is subject to certain restrictions, and it can take some time when creating large raster datasets. Whether a raster dataset consists of one or a mosaic of many input rasters, raster datasets are converted to ERDAS IMAGINE (IMG) files and stored in a special folder at the same level as the personal geodatabase (MBD) file. This folder and its contents are managed by the personal geodatabase.

Rasters in catalogs cannot be used for analysis as a single combined entity, but individual raster dataset items within raster catalogs can. Raster catalogs are very useful for displaying data and managing archives even when the rasters partially or fully overlap.

When a raster catalog is the choice, decide whether the catalog will be managed or unmanaged. An unmanaged catalog is a collection of rasters that may be in different formats with different cell sizes, pixel depths, and other characteristics. These catalog raster items may even have different defined spatial references. Raster data is referenced through the RASTER field in the raster catalog table, and the path name is used to find the raster dataset file on disk. In a managed catalog, rasters are copied to IMG files, which are stored in a special folder with the same name, and on the same level, as the personal geodatabase. Whether nonmanaged or managed, a new catalog is created within the geodatabase and multiple rasters are added to it.

Conclusion

By managing rasters in a personal geodatabase, all data—vector and raster—can be viewed from a central location. There are many powerful tools for managing raster data within the personal geodatabase. Additional information

on working with raster data in ArcGIS is available online from the ESRI Virtual Campus. *Working with Rasters in ArcGIS 9* is a course that teaches students how to use core ArcGIS tools to work with raster data. *Georeferencing Rasters in ArcGIS*, a workshop, focuses on the tools in ArcGIS 9 used to georeference, transform, and project raster data. Participants will learn how to transform a raster to fit a known projection including scaling, shifting, rotating, and projecting raster data.

About the Author

Chris Wayne is the principal at Mazama Mapping, a GIS consultancy specializing in field GIS applications, GIS services, and GIS training. He is an Authorized ESRI Instructor certified to teach three courses and also develops custom coursework. He worked for ESRI for more than three years and specialized in ArcPad, ArcGIS, and ArcGIS analysis extensions. He can be reached at mazama_mapping@hotmail.com.

Continued on page 56

Managing Rasters in a Personal Geodatabase

Continued from page 55

Personal Geodatabase Storage Type	Input Raster Format	Personal Geodatabase Raster Format	Path to Raster Files	Used for Analysis?
Raster Dataset (single input , always managed)	Any file-based ArcGIS supported raster data	Single copied IMG file	Subfolders within workspace/ <gdbname>.idb folder	Yes
Raster Dataset (multiple inputs , always managed)	Any file-based ArcGIS supported raster data	Single IMG file; spatial reference conversion and cell-size resampling upon input	Subfolders within workspace/ <gdbname>.idb folder	Yes
Managed Raster Catalog	Any file-based ArcGIS supported raster data	Multiple copied IMG files	Subfolders within workspace/ <gdbname>.idb folder	Individual raster items can, but not entire raster catalog as a single entity
Nonmanaged Raster Catalog	Any file-based ArcGIS supported raster data	Multiple raster files of any supported format	Uses source data directly	Individual raster items can, but not entire raster catalog as a single entity
Managed Rasters as Attributes (of features: points, lines, and polygons)	Any file-based ArcGIS supported raster data	Multiple copied IMG files	Subfolders within workspace/ <gdbname>.idb folder	Rarely used for anything other than viewing
Nonmanaged Rasters as Attributes (of features: points, lines, and polygons)	Any file-based ArcGIS supported raster data	Multiple raster files of any supported format	Uses source data directly	Rarely used for anything other than viewing

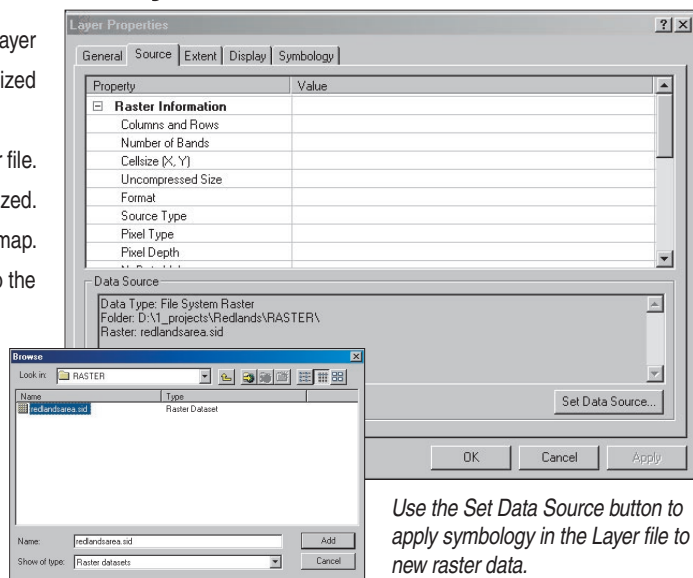
Comparing raster storage methods

Reuse Symbology on a Raster Layer

It is easy to reuse symbology created for one raster layer with another layer by creating a Layer file and redirecting the data source for the symbolized Layer file.

In ArcMap, save the raster layer with the desired symbology as a Layer file. Add the .lyr file to the map containing the raster that needs to be symbolized. Open the Layer Properties dialog box for the Layer file just added to the map. Click on the Source tab and click the Set Data Source button. Navigate to the raster data that needs to be symbolized and choose it.

Alternatively, a DLL plug-in created by Trent Hare is available online from the ESRI Support Web site. It can copy an image's symbology and paste it to other images or image catalogs. This was written to easily set multiple image catalogs or many images to have a NULL background. Visit support.esri.com and search for "Copy and Paste Image Symbology (w/Nulls) to Images and/or ImageCatalogs."



Use the Set Data Source button to apply symbology in the Layer file to new raster data.