Comparison of managed raster datasets (ArcSDE technology) and the ArcGIS Server – Image extension (ArcGIS Image Server technology)

ArcSDE's raster data management capabilities and the Image extension for ArcGIS Server are complimentary products. ArcSDE is used to store pre-processed raster data within a DBMS and then serve the data to multiple users and applications in a variety of environments. ArcSDE leverages the robust nature and capabilities of the geodatabase's underlying DBMS. The Image extension to ArcGIS Server is a server technology that in response to a client request quickly reads, mosaics, and optionally processes the raster datasets that are then returned to the client application as a single image. The Image extension performs processing of imagery on-the-fly and this can be used to create multiple image services from a single source. The data can be stored as files or within ArcSDE.

The differences between these two products are not specifically in what they do, but what they aim to do. The ArcSDE technology is primarily for DBMS-based data storage and access, and returns tiles of rasters stored in the ArcSDE geodatabase to client applications. The Image extension enables large sets of raster data (imagery), stored in many different formats, to be cataloged, then served to client applications as mosaicked image services with image processing performed on the server.

ArcGIS Server

Image services were added to ArcGIS Server at version 9.3. Image services are a new service type optimized to serve imagery to Web and desktop applications. Similar to map and globe services, Web applications make request to image services using a number of different protocols and standards such as WMS, WCS, SOAP, and REST. The server reads the required imagery and returns an image, with the required processing, to the client application. Depending on the protocol used, the client can define parameters for the required imagery such as the projection to be used, sampling method, as well as the format and compression for transmission. This enables client applications to quickly obtain the imagery required for display or analysis as well as work over low bandwidth networks. A significant advantage of image services is that they not only return imagery as 'pictures' to the client application, but can also return the pixel values required for use in analysis or as elevation models in globes.

The source to an image services can be any raster dataset such as a TIF, NITF, MrSID or a raster dataset from ArcSDE. If the Image extension to ArcGIS Server is licensed, then the source can be a compiled image service definition (ISCDef). An ISCDef is created using the Image Service Definition Editor within ArcMap and defines a collection of raster datasets along with the definition of processing parameters and metadata for each raster.

When ArcGIS Server receives a request for imagery, if the source is a raster dataset then the server extracts the required imagery from the source, reprojects and optionally compresses the data for transmission. For image services whose source is the Image extension, the requests for processing is passed to the Image extension to perform the required mosaicking and processing on-the-fly, returning an image that is passed back to the client application.

Managed raster datasets (ArcSDE technology)

Organizations using ArcSDE technology for raster data management have relatively stable collections of preprocessed imagery (such as a series of DOQQs that get flown every 3 years) that they want stored with their vector data in a DBMS-based ArcSDE geodatabase. These organizations want preprocessed data, so their users have fast and easy access to the raster data over both internal networks and over the internet without the need for additional processing. Access is either through ArcGIS Desktop products, or with ArcIMS or ArcGIS Server services.

Key advantages of the ArcSDE Raster format

- Central storage of the raster data in a DBMS resulting in the total integration of an organization's spatial data (imagery, vector, metadata, address, etc.) in the same framework
- Security and access is controlled by a DBMS and via ArcGIS tools
- Administrators can use standard database tools such as backup, recovery, replication, for central administration
- Central storage provides a central data hub for all clients to load data into and access
- Enables a user to maintain metadata with each raster dataset or collection of raster datasets within the database
- Optimized for fast, seamless, distribution to clients, as the raster data is preprocessed into tiles in the DBMS
- Data stored in ArcSDE can be utilized simultaneously in a multi-user environment throughout the organization by desktop applications, Internet mapping servers, and custom applications built with ArcGIS Engine
- Raster datasets are full 'citizens' of the geodatabase (therefore able to take advantage of spatial searching)
- Good raster data performance to all users through indexing, tiling, pyramids, and definable compression techniques
- Cross platform (Multiple operating systems, such as Windows, Linux, and Sun Solaris and multiple DBMSs, such as IBM's DB2 and Informix, Microsoft's SQL Server, and Oracle)

Image extension to ArcGIS Server

Organizations wanting to use the Image extension to ArcGIS Server may have small to vast collections of existing, dynamic, or newly acquired raster data (imagery) in either a preprocessed or non-processed form. The data may exist in many formats, such as TIFF, JPEG 2000, and geodatabase rasters, as well as specific vendor outputs, such as QuickBird, USGS DEM, and MATCH-AT. Organizations may have a need for: rapid access to huge raster datasets after acquisition; for a large number of users to access the data from different applications; to process and distribute huge quantities of imagery (such as a data provisioning company); or to archive and easily access large quantities of raster data in both the original raw format and in a processed format. Access is either through ArcGIS Desktop or Server products or through client 'plug-ins' for non-ESRI products such as AutoCAD® and MicroStation®.

The Image extension combines together the traditional image processing and distribution workflows—enabling administrators of raster data to maintain only the primary raster data, with multiple image products created instantly, on-the-fly, as required by users. Whereas, the traditionally separate workflows cause data redundancy, can explode into volumes of data, and may hinder efficient data management. Improving the traditional workflows for raster data management and distribution, results in improved efficiency and reduced costs.

Key advantages of the Image extension

- Reduced delays from data (image) acquisition to dissemination
- Dynamic mosaicking and processing of multiple images
- Simplified data management
- Elimination of data storage at different stages of processing by processing on-the-fly, thereby, reducing the data duplication
- No loading is required, thereby making mosaicked raster datasets available as an image service more quickly than with traditional workflows (Note: overviews may need to be created, but they are a fraction of the data volume)
- Standard file management tools can be used for backup, recovery, and replication
- Support of all ESRI raster formats
- Multiple image services can be produced from the same source data, such as georeferenced or orthorectified imagery, using different band combinations or image fusion techniques (e.g., pan-sharpening), and applying various enhancements
- Higher quality image output as there is a single sub-sampling from raw data to the display
- Superior handing and easy access to overlapping raster datasets
- Client access to metadata about the individual rasters making up a view

Considerations if choosing from these two raster solutions

In most cases the two technologies will work well together as a raster solution for your organization. However, you may want to make a decision to use one of the two, or if you are using both you may be deciding which product would provide the best solutions for a particular need. These decisions require many considerations which are discussed below.

Platform support

The ArcSDE technology is supported on a number of platforms including: Windows 2000 Server, Windows Server 2003, Linux-Intel Red Hat Enterprise 3.0 AS/ES, Linux-Intel SuSe Enterprise Server 9, IBM AIX 5.x, HP-UX 11i, HP Tru64 5.1a and 5.1b, and Sun Solaris 8, 9, and 10. (Note: Not all databases are available on all operating system environments.) See, ArcSDE System Requirements for the most current list.

(http://support.esri.com/index.cfm?fa=knowledgebase.systemRequirements.matrix&pName=ArcSDE&ProductID=19&pvName=9.0%20Pre-release&PID=19&versionID=38&PVID=51)

The Image extension to ArcGIS Server is available on Microsoft Windows XP, VISTA Enterprise, and Server 2003 & 2008. However image services can work from ArcGIS Server on Unix and Linux. The images stored can be on any accessible file system or from ArcSDE.

Formats

Raster data is stored (converted) using the ArcSDE technology as an ArcSDE raster as either a raster dataset or a raster catalog. Ancillary files can be ingested and organized along with the raster data. Whereas when using the Image extension the data storage remains in a file system or in a DBMS. By keeping the data stored in a file system all the associated files can be kept together, such as meta information, as it is provided from the vendor. The image service definition is a data model created using the authoring tools for the Image extension which creates a series of files organized on the file system. These image service definitions are compact because they do not contain any pixel data and they can easily be moved to other locations. The raster data is linked to (or indexed) within this definition—similarly to an unmanaged raster catalog in ArcGIS.

Image processing

When using the ArcSDE technology, raster datasets may be mosaicked into a single raster dataset when the data is loaded into the DBMS. If there is any overlapping raster data, the overlapped data is lost. With the Image extension the raster datasets are stored as individual rasters and are dynamically mosaicked together on demand. You also have the ability to edit a polygon that defines the viewable data of each raster dataset, therefore, you do not need to clip and loose any source data or have to make a copy. To save all the data using the ArcSDE technology, you would have to create a raster catalog which stores and displays each separate image.

Additionally, other image processing functions can be applied to the raster data on demand, with the Image extension. These processes include enhancing the appearance; applying algebra to the pixels and bands; changing the output colors; visualization of elevation data; and applying geometric transformations such as reprojecting, warping, and orthorectifying. Many different methods for mosaicking are allowed, enabling the overlap of imagery to be exploited. Processing is performed on the server so it has no load on the client application.

At 9.3 the image services can serve raw data values (as a raster dataset), therefore, processing can take place on the client machine, using the image service as an input raster source. Therefore, raster data served via the ArcSDE technology or from the Image extension can be used as input for geoprocessing, or if a DEM, as a source for an ArcGlobe surface.

Security

At 9.3 access to imagery from image services is controlled via the Web tier through ArcGIS Server. Both the ArcSDE technology and Image extension allow direct connections to the services to be made to access data without going through the Web tier. Such direct connections use protocols, such as RPC. For the ArcSDE technology, the security of such connections is controlled by the underlying DBMS security model. With the Image extension direct connections can be set up using access restrictions that are based on Windows domain authentication.

Image update cycle

With ArcSDE rasters, new raster datasets can be loaded at any time and existing data can be updated. However, the data must be updated in the server and large volumes of data updates may take some time.

With the Image extension raster data can quickly be made available to end users with no, or minimal, data preprocessing. Published image services can be updated and altered without affecting the end user. Therefore new data can be added, data can be replaced, or properties of the image service can be modified without bringing down the image service.

Client-side access

At 9.3 most client applications obtain access to imagery by accessing image services. Through image services a wide range of both desktop and Web applications can quickly access imagery using protocols such as SOAP, REST, WMS, or WCS. Legacy applications can also make direct request for data from both the ArcSDE technology and the Image extension.

Using direct connections, the ArcSDE technology supports a wide variety of data access, including the ability for ESRI's ArcGIS product line and developers (via a C and Java API) to query, display and extract raster data, and access to raster and vector data. The client application request the individual tiles of raster data from the server and then, performs the limit processing to stitched together, re-project and display them as required. When accessed via ArcGIS Server this processing, and the optional compression for transmission, is performed by the server.

The Image extension has a direct connection client that is included in ArcGIS product line as well as existing plug-ins for a number of non-ESRI products including MicroStation and AutoCAD. When a client application makes a request for data from an image service, the coordinates of the current view along with some properties are sent to the server which processes and returns the required image. For many client applications the required imagery is pasted to the application's display and requires very little processing or memory by the client application. Additionally, new client applications can be developed. The client application can modify some service properties such as the mosaicking method, sampling method, and the compression for transmission method. Imagery can be extracted from the client and stored to a file within limits defined by the administrator. When accessed as image services the client requests are passed through ArcGIS Server to the Image extension with very little load placed on ArcGIS Server.

Speed and scalability

In most cases, both speed and scalability between the products are similar and tend to be dependent on hardware and data. Due to the large volumes of data, it is not possible for the server to maintain all imagery in memory so the data must be read from disk. The key aspect affecting performance is the format of the data on disk. When using ArcSDE rasters the processing of adding the data ensures that the data is structured as a quickly accessible tiled structure with optional compression. These tiles are read from the disk by the underlying database.

When using image services the pixels accessed can come from the same database tiles or directly form the original source files. Most image file formats are structured for fast access and can be accessed as fast as through a database. There are some file format that are advantageous to first convert.

Both technologies are scalable. The inherent scalability of ArcSDE rasters come through the scalability of the underlying database server. The ArcGIS Image Server technology has specifically been designed to quickly read and processes imagery, use multiple threads and cores available on servers, and can be run concurrently on multiple server machines. There are three stages in the data access and processing that can affect speed: reading the data, processing the data, and transmitting the data. The server must first read the pixels, which is optimized, but the performance can be degraded when using compression formats that are CPU intensive to decompress (such as a wavelet compression). The on-the-fly processing is very efficient and most processes provide little additional demand on the server. Image services that use several computationally intensive processes, such as orthorectification and pan-sharpening, can slow the performance by a factor of 2 to 3 times. The compression of the data for transmission only puts a small load on the server, but can substantially speed up the transmission of imagery over lower bandwidth networks.

Optimum scalability for serving imagery to the Web is gained by using map caches. Creating a map cache results in pre-processing the imagery into small tiles that are directly accessed by Web applications with nearly no CPU load on a server. Creating a cache is similar to creating ArcSDE rasters, but the tiled images are stored as individual tiles on a file server and security is not managed by the database. Caches have similar disadvantages to mosaicked raster datasets in that information from overlapping images and metadata is lost. The compression used for transmission is the same as the compression used to store the data; therefore, it cannot be changed later. This pre-mosaicking of imagery often results in a loss of image quality due to multiple sampling of the imagery. Creating a cache is the most scalable technology for serving imagery. ArcGIS Server enables the creation of map caches from any map document which can include raster datasets directly from files, ArcSDE or the Image extension. Generating a cache can be time consuming; however; ArcGIS Server map caches can be generated on demand. This enables caches to be defined for very large areas without the need to pre-process the complete area.

Overlapping images

An image service definition generally contains multiple raster datasets. And like the raster catalog these datasets can overlap and have different pixel resolutions and extents. When accessing the raster catalog the user will view the same raster dataset or datasets irrespective of the display resolution. However, in an image service, multiple image sources with multiple resolutions will display according to their resolution. For example, if you have scanned map sheets including 1:5000, 1:12,000, 1:24,000, and 1:100,000. When you first see the image service you are likely viewing the entire service and likely at a resolution at or above 1:100,000, therefore you will view the 1:100,000 scale map sheets. When you zoom in you will eventually see the 1:24:000 once you reach that scale, and eventually the 1:5000 when you've zoomed in to the highest resolution. However, you still have access to the lower resolution map sheets if you choose to view them.

When accessing multiple overlapping images the methods by which imagery is returned differs. When using the ArcSDE technology, raster datasets may be mosaicked into a single raster dataset or loaded separately into a raster catalog (they can also be loaded as separate raster datasets). If there is any overlapping raster data loaded into the mosaicked raster dataset, the overlapped data is lost. When accessing a raster catalog, the portions of every raster dataset (even overlapping) covering the display are transmitted from the server to client (which can be slow if there are many overlapping raster datasets). When using the Image extension, the client defines the order of the overlapping imagery and the server performs the mosaicking. Therefore, the overlapped

information is always accessible, but only transmitted when requested, to reduce the amount of imagery transmitted. If required, the individual overlapping images can be accessed and used, for example with the Swipe tool for change detection.

Metadata

Similar to how overlapping imagery is handled, metadata for individual images is lost when a mosaic is pre-created, but is available in a raster catalog by querying the database. With the Image extension the metadata is maintained and is available to the client applications.

How do they complement each other?

Using the Image extension and storing raster datasets in ArcSDE generally compliment or supplement each other. Both technologies are important components of a complete image management solution. The Image extension can use ArcSDE raster datasets as an image format allowing dynamic processing of this data to be served to other users. In addition, the Image extension can also be used to prepare raster data for loading into ArcSDE or creating map caches.