

Enterprise Image Management

An Esri® White Paper
November 2015



380 New York Street
Redlands, California 92373-8100 USA

800 447 9778
info@esri.com

esri.com

Copyright © 2015 Esri
All rights reserved.
Printed in the United States of America.

The information contained in this document is the exclusive property of Esri. This work is protected under United States copyright law and other international copyright treaties and conventions. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, except as expressly permitted in writing by Esri. All requests should be sent to Attention: Contracts and Legal Services Manager, Esri, 380 New York Street, Redlands, CA 92373-8100 USA.

The information contained in this document is subject to change without notice.

Esri, the Esri globe logo, ArcGIS, esri.com, and @esri.com are trademarks, service marks, or registered marks of Esri in the United States, the European Community, or certain other jurisdictions. Other companies and products or services mentioned herein may be trademarks, service marks, or registered marks of their respective mark owners.

Enterprise Image Management

An Esri White Paper

Contents	Page
Overview.....	1
Enterprise Requirements.....	1
ArcGIS as a Comprehensive Imagery System.....	2
Using Imagery.....	2
Ready to Use or Manage Your Own.....	3
Image Access	4
Imagery Management and Processing	5
The Mosaic Dataset.....	5
Formats and Compression Support.....	5
Adding Data to a Mosaic Dataset	6
On-the-Fly Processing	6
Handling Data Extents and Nodata.....	7
Dynamic Mosaicking.....	7
Pyramids and Overviews	7
Refining Processing Parameters	7
Mosaic Dataset Properties.....	8
Scaling to Massive Collections.....	8
Automation of Workflows	10
Scaling to the Cloud.....	11
Optimal Data Transfer	11
Open Standards	12
Client-Defined Processing	12
Creating and Serving Tile Caches.....	13
Geoprocessing.....	13
Raster Analytics and Big Raster Processing.....	13
Conclusion	14
Getting Started	14

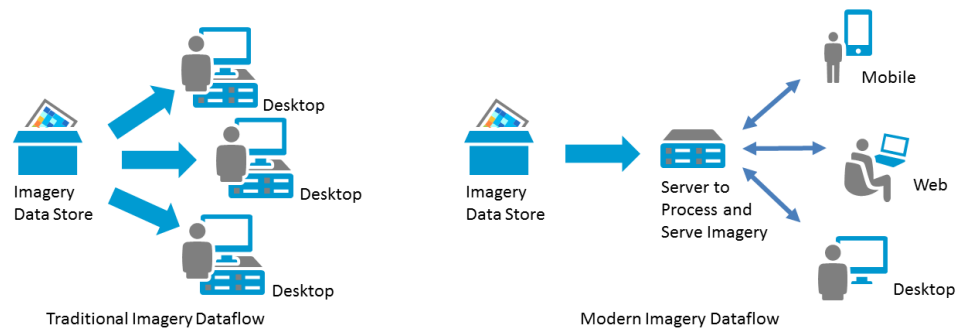
Enterprise Image Management

Overview Massive volumes of imagery and remote-sensing data exist. These datasets have significant value when the information inside them becomes accessible. ArcGIS® is a commercial off-the-shelf (COTS) platform that provides extensive enterprise image management capabilities and is used by organizations in a wide range of industries to manage their imagery holdings, making them accessible and turning them into useful information products for both visualization and analysis.

Enterprise Requirements Organizations ranging from state and local government to natural resources, utility, defense, and emergency response all collect and maintain large collections of imagery and remotely sensed data. Traditionally, there have been challenges in managing such data and making it accessible to a wide range of users in different applications. These challenges have been accentuated in recent years by the ever-increasing volume and diversity of remotely sensed data. The data comes from a wide range of sensors on platforms ranging from satellites, aircraft, and drones to terrestrial vehicles. Much of this data may exist in the form of raw pixels directly from the sensors, or it may have been preprocessed into finished information products such as orthophotos, elevation models, or thematic maps. Collections of such images also come from scanned maps, aerial films, and historical photo prints. Sometimes these datasets, also referred to as "rasters," may be the output of image processing algorithms or the result of some type of spatial analysis workflows. These datasets are frequently the most accurate and up-to-date depiction of actual features and conditions on the ground for any given area of geography and time. This information is of high importance to the organizations, as it is used to make important decisions.

The quantity and use of imagery to be accessed are varied. Some users need to access individual datasets, for example, to perform analysis on a specific satellite scene. Others need to access extensive collections of images as a single product, such as a mosaic from multiple orthophotos or a single elevation surface created from multiple sources. In many cases, the dataset may represent multiple temporal images from sensors or a multidimensional dataset of scientific data, and the users need to obtain temporal profiles or quickly view or analyze a combination of the measurements.

The massive volume and sheer size of imagery data makes it impractical to move the data to the application for processing. A more practical and efficient approach is to serve the imagery as web services through servers that process the imagery dynamically and only transmit the information requested to the application. In this way, the imagery data is instantly available, and the time involvement and hardware costs are significantly reduced.



ArcGIS as a Comprehensive Imagery System

ArcGIS provides enterprise image management as a component of a complete geospatial platform. The ArcGIS platform enables organizations to fully exploit and leverage geospatial information as an integral part of the organization. As a comprehensive imagery system, ArcGIS includes an extensive range of image management, processing, and analysis capabilities used to serve imagery efficiently from both enterprise or cloud infrastructures.

Using Imagery

Imagery on its own provides useful context, but the true and enhanced value comes from visualizing and analyzing imagery georeferenced with other datasets. In its simplest form, natural color imagery is used within a geographic information system (GIS) as a backdrop to provide context for overlaying a map and allowing users to visually interpret what is on the ground or investigate change and make updates to the map. The value is increased when the full information content of the imagery becomes accessible to users in software applications for advanced visualization and analysis, enabling the extraction of information to form understanding and knowledge. These applications can range from focused apps running on mobile devices to generic web applications or desktop GIS workstations optimized for geospatial analysis. The numbers of users accessing the imagery may range from a few individuals to very large numbers of users in publicly accessible applications.

The ArcGIS platform provides an environment and tools to create geospatial applications that can easily consume and provide access to imagery for both visualization and analysis. Focused, feature-rich web, mobile, and desktop applications that enable the integration of imagery and other map data can be quickly created using templates. ArcGIS also provides developer application programming interfaces (APIs) and software development kits (SDKs) that enable users to build robust, customized applications. These applications can make requests to servers, which in turn perform the processing dynamically and return the results quickly. Web applications can also request the streaming of data values from the server and perform additional client-side processing to create highly interactive applications.

The ArcGIS APIs enable extensive filters and queries to be set on the image collections. This provides users with access to massive collections of overlapping imagery, such as temporal collections of multispectral satellite imagery, yet work with only the imagery that fits their individual needs. The applications can also interact with geoprocessing tasks running on the servers that can utilize the full set of ArcGIS tools to perform analysis on collections of raster, vector, and tabular geospatial data and then return the results possibly in the form of new maps or features.

The same server-based capabilities are also accessible to desktop applications, such as ArcGIS for Desktop. The desktop provides a more generic visualization and analytics environment to work with all forms of geospatial data, including imagery, which is accessed directly from local storage or through the same web services used by the web applications.

Native applications running on mobile devices such as Android, iOS, or Windows phones can be created using ArcGIS Runtime. These applications can either access imagery locally on the device or connect to servers to access the imagery. An application for field data collection may use imagery only as a backdrop. A more advanced application may use raster data, such as terrain elevation, for analysis—such as performing line-of-sight computations for military situational awareness.



Here is an example of a web page (esri.com/landsatonews) with an embedded application that provides access to hundreds of terabytes of temporal, multispectral Landsat imagery.

Ready to Use or Manage Your Own

ArcGIS provides ready-to-use content as well as the tools required to manage and serve your own collections of imagery. ArcGIS provides access to large collections of online imagery via the Living Atlas of the World (livingatlas.arcgis.com). This includes a World Imagery Basemap that provides global 1-meter or better resolution imagery, typically two to five years old, with new and high-resolution imagery available for the US and other parts of North America. Global datasets curated by Esri and made accessible directly as web services include Landsat 8 and GLS, GLDAS, NAIP, MODIS, and elevation. Other organizations also publish publicly accessible datasets via ArcGISSM Online. Additionally, users can subscribe to imagery from a number of satellite and aerial imagery providers. Access to imagery through these online services enables organizations to immediately start using and derive value from these multiple data sources.

Many organizations have also amassed large collections of imagery and rasters. ArcGIS provides the platform to transform these often hidden resources into valuable information products. With ArcGIS, the imagery and raster data can be cataloged into mosaic datasets and then served as image services that make them accessible to a wide range of applications. ArcGIS as an enterprise image management system provides all the tools and technology to manage and serve your imagery and raster data. The services can be used internally within an organization, shared securely externally, or made publicly accessible.

Image Access

Having access to imagery quickly when it's needed increases its value and is a key requirement in organizations. ArcGIS supports four different image access methodologies; each has advantages and disadvantages:

- **Download**—The required data is downloaded from servers to local storage for use in the application. This traditional clip-zip-ship type of image access has traditionally been used for disconnected applications but is typically avoided, as it results in potentially significant download times and requires additional storage, data management, and processing by end users. The primary use case for such access is when a user needs to take an extract of the data to a location where there is insufficient Internet access to a server. The data they extract may be just a set of pregenerated tiles as a background or the full dataset for use in analysis.
- **Tile cache**—This is used primarily to provide background imagery in many applications. For tile cache, the imagery is preprocessed into a highly compressed tiled format, stored on servers, and streamed on demand to the applications. Tile cache is the most efficient way of serving 8-bit RGB background imagery to applications, as it requires minimal load on servers and maximizes the use of caching on both server and client. Such tile cache provides good background imagery but limited analysis capabilities. The primary use case is as background imagery, such as the World Imagery service on ArcGIS Online, or similar background imagery services created from smaller datasets.
- **Dynamic image services**—Client applications interact with servers that have access to the imagery. Requests are made to the servers, which quickly access the data, process it into the required product, and return the results. Image services enable applications to access massive collections of minimally processed imagery, with the server applying processes, such as image enhancement, computation of a vegetation index, or image classification, but only to the imagery within the area requested by the client. These dynamic image services provide access to the full information content of imagery. The requests are synchronous with the image being returned nearly instantaneously. ArcGIS provides a very rich REST-based web service protocol that enables the client to define the processing to be applied on the server.
- **Geoprocessing services**—The client makes requests to a task running on the server that can process all forms of geospatial data. Geoprocessing services can access and process many different raster, vector, or tabular datasets with the result often being the answer to a specific question or a new set of features. These services are primarily used when the output is not an image. A typical example of geoprocessing from rasters is the computation of viewsheds from elevation data, where the source

point may be outside the extent of the current view, so it cannot be computed based on the viewable imagery.

Imagery Management and Processing

Managing imagery and raster data efficiently and correctly is key to ensuring accessibility. For this, ArcGIS uses a geodatabase structure called **mosaic datasets** to manage the imagery. In addition, ArcGIS includes all the technology and the associated tools needed to build and maintain mosaic datasets.

The Mosaic Dataset

A mosaic dataset is a well-defined geodatabase structure optimized for working with large collections of imagery and rasters. Mosaic datasets are stored as either a file geodatabase or an enterprise geodatabase such as Microsoft SQL, Oracle, or PostgreSQL. The imagery and raster data does not need to reside in the database, and most organizations instead store it as files on enterprise or cloud storage. A single mosaic dataset can reference millions of images and make them appear as a single virtual dataset or enable quick access to any individual image or collection of images.

With mosaic datasets, the large volume of pixel data (contained in the imagery and rasters) is not loaded into the database and is instead referenced. The metadata about the data sources, as well as information on how to process the imagery into different products, is stored in the mosaic dataset. When a request for imagery is made, the mosaic dataset is used to determine what images are required and what processing is to be applied. Only the required imagery is read, processed, and returned.

A mosaic dataset can be authored directly in ArcGIS for Desktop and ArcGIS Pro (Standard) using a wide range of tools. In many organizations, the process of authoring mosaic datasets is automated by using a set of Python-based templates that automate the full process. The mosaic datasets can be used directly in all ArcGIS for Desktop applications. When accessed in this way, the required data access and processing is performed by the desktop application. Alternatively, the mosaic dataset can be served as image services through ArcGIS for Server and then accessed through the very wide range of applications. Image services enable client applications to make requests to the server, and the server performs the required data access and processing, returning only the required pixels to the client application.

Formats and Compression Support

ArcGIS can read imagery in nearly every known raster format, so the data can be left in its original form, removing the need to make further copies and so reducing storage requirements. Typically, the source imagery is stored uncompressed or with lossless compression so that there is no information lost. Many organizations also receive or generate preorthorectified and mosaicked imagery, which often have lossy compression applied. The lossy compression can be used to significantly reduce the data volume, but this also reduces the ability to perform analysis on the data.

There are many different image formats; some are better suited than others for fast access. Not all formats are optimum for data access, and it is sometimes advantageous to convert the data into an optimum format while still maintaining the structure and metadata of the original data. Conversion to a new format can also provide the opportunity to compress the files and so reduce storage and improve access performance. ArcGIS supports a wide range of lossless and lossy compressions including LZW, JPEG, and wavelet-based compression such as JPEG2000. ArcGIS also supports Limited Error Raster Compression (LERC), a "controlled lossy" compression that enables fast

compression and decompression of imagery while maintaining specified tolerances. This is of significant value when working with elevation and scientific data that cannot be compressed using traditional lossy compression methods. LERC is also optimum for lossless compression due to its fast speed and the good lossless compression ratios it is able to achieve. LERC is therefore recommended for the storage of elevation data and high bit depth satellite scenes that are typically stored using lossless compression, as it reduces the data storage volumes and increases access performance.

Adding Data to a Mosaic Dataset

Many raster datasets only have minimal metadata, such as georeferencing, number of bands, and bit depth. These raster datasets can be added to a mosaic dataset using the Add Rasters to Mosaic Dataset tool that will crawl through directories and ingest all the relevant metadata. ArcGIS also supports a wide range of raster products from imagery providers that contain sensor-specific metadata. These are added to mosaic datasets using the same tool but using raster types that are configured for the specific sensors. The directories are then crawled to extract all the detailed metadata from the vendor products. The raster types can also set up appropriate functions, such as orthorectification and pan sharpening, which transform the imagery into higher-value products. The use of raster types further simplifies the ingestion process.

Many organizations also have existing databases or tables that provide metadata and other parameters about the datasets. Such tables can be added to the mosaic dataset using the table raster type to incorporate the associated metadata. At any time, additional metadata and attributes can be assigned to each raster in the mosaic dataset. The metadata ingested is used to aid in defining the correct processing as well as to enable queries and filters.

On-the-Fly Processing

A range of functions can be associated with each raster in a mosaic dataset to transform the pixels from data values stored on disk to values required by the applications. These functions are applied as on-the-fly processes when the rasters are accessed.

There is a wide range of both geometric and radiometric functions that can be applied. The geometric functions transform the location of the pixels and are used to perform processing such as reprojection of imagery between different coordinate systems or georeferencing such as orthorectification of satellite or aerial imagery. The radiometric functions transform the pixel values and are used for processes such as image enhancement, pan sharpening of satellite imagery, or the computation of vegetation indices and band arithmetic. These functions are able to work with the available bit depth of the imagery, enabling the full dynamic range of the imagery to be exploited. This is especially important when working with newer satellite and aerial imagery or with elevation and scientific datasets.

ArcGIS provides a large library of these functions, which can be chained together to perform a wide range of data transformations. For users requiring additional raster analysis, the Python adapter function can be used. This enables Python-based functions to be developed that utilize the extensive libraries of both NumPy and SciPy. Python raster functions are optimum for the development of more advanced algorithms such as change detection or the processing of scientific multidimensional data.

Handling Data Extents and Nodata

Rasters are by definition rectangular, but may contain areas of no data (usually on the edges) where valid pixel values do not exist. ArcGIS provides extensive support for defining the extents and handling of nodata areas using a combination of footprints and nodata masks. Each raster in a mosaic dataset has an associated footprint geometry that is used to quickly search for rasters covering a specified extent but can also be used to virtually clip each raster. Footprints are computed when the rasters are added to a mosaic dataset but can be further modified to exclude pixels from processing or display. Similarly, nodata pixel masks can be used to exclude pixels. Using footprints and nodata masks provides flexibility and optimization, especially when working with collections of overlapping imagery.

Dynamic Mosaicking

Since the mosaic dataset is based on databases that contain the metadata, it is simple to define queries to return details of imagery covering a specified area of interest. Users can also define filters in the form of Where clauses so as to display or work only with subsets of data.

For any specific extent, there are likely to be a number of overlapping images. Dynamic mosaicking is the capability to define and refine the order in which images should be merged or blended. Each mosaic dataset has a default mosaic method that defines the default rules for ordering overlapping imagery. There are a range of these mosaic methods, and they can be changed by the users. The ByAttribute mosaic method enables a user to order the imagery based on an attribute, such as date or cloud cover, so that the latest or least cloudy images appear on top. The closest to nadir method is used often for collections of aerial imagery so that the most top-down-looking image is displayed on top. The seamline mosaic method can be used to create seamless mosaics by which the images are blended along predefined seamlines. Specific fields in the mosaic dataset table can also be used to force individual images to be displayed first or last. The mosaic method and query filters provide the ability for users to work efficiently with overlapping images.

Pyramids and Overviews

Mosaic datasets contain large collections of imagery, yet users still expect to be able to zoom out to see the collections at a small scale over a wide area. Opening and processing thousands of images for each request is not scalable, so the mosaic dataset makes use of two different forms of reduced resolution imagery: image pyramids and service overviews. Pyramids are associated with each source image and provide fast access if the image is to be visualized at a lower resolution. Similarly, at the mosaic level, the mosaic dataset has a set of overviews, which are compiled from source images and are displayed when a user zooms out to a small scale. Such overviews are small in volume when compared with the source data. In this way, the system provides fast access to the imagery at all scales.

Refining Processing Parameters

The mosaic dataset stores metadata and attributes about the rasters and the processing chains to be applied. These processing chains can apply a wide range of geometric and radiometric processes on the imagery. ArcGIS provides a range of tools to set and refine these parameters. These tools include the following:

- **Build Footprints**—Used to recompute the footprint of an image, for example, to clip off unwanted areas such as large nodata areas or the collars from scanned maps

- **Edit Raster Function**—Used to add new functions, such as image enhancement, to a selection of images, or to change the parameters of existing functions. (For example, the accurate orthorectification of imagery requires appropriate sensor model parameters and elevation models. These parameters can be set during ingestion as part of the raster type or later refined using Edit Raster functions.)
- **Image Registration**—Manual and automated image registration that can be used to refine the geometric transforms to be applied to the imagery so that they better fit the ground (The automated image registration tools will automatically detect control points based on existing imagery such as the World Imagery basemaps or other imagery data sources.)
- **Block Adjustment**—For satellite imagery, used to collect tie points between overlapping images and performing a block adjustment along with control points so as to minimize the displacement between all images
- **Seamline Generation**—Used to determine the best locations for seamlines that define where to blend images together
- **Color Correction**—Determines adaptive enhancements to be applied to each image so as to minimize the radiometric differences between the images

The parameters for such processes can also be ingested from other image processing packages, making it possible to integrate processing parameters from existing production systems, as well as enable a number of Esri partners to integrate their specialized image processing capabilities with ArcGIS.

Mosaic Dataset Properties

In addition to attributes and properties specific to each raster, the mosaic dataset also defines a set of properties that are applicable to all the rasters. Some of these mosaic dataset properties are used to define defaults for when the imagery is first accessed, for example, the default mosaic method for blending the imagery or the compression to be applied prior to transmitting the imagery to the client application. The mosaic dataset properties are also used to define a range of functions that can easily be applied to the imagery. For example, a mosaic dataset of elevation data may have a set of predefined functions for hillshade, slope, and aspect maps that are computed on the fly as the data is accessed.

Scaling to Massive Collections

Mosaic datasets are scalable. They can be used to work with a few images or many millions of images. The images can be added, edited, or removed at any time, enabling the implementation of mosaic datasets for dynamic environments where new imagery is frequently received.

Mosaic datasets work with a wide range of imagery. A simple example is the preprocessed orthophotos that state and local governments have. Large collections exist from multiple years, often with different resolutions and projections. Some older imagery may be panchromatic, while others are natural color or include an infrared band. Using a mosaic dataset, all the imagery can be managed either as separate mosaic datasets per collection or combined into a single mosaic dataset, then seamlessly accessed without duplication or reprocessing. An example of such a service is the NAIP services on

ArcGIS Online that provides access to about 100 TB of current and historical NAIP 1-meter 4-band imagery for the United States.

The source and resolution of the imagery in a mosaic dataset may also be varied. For example, mosaic datasets can be used to work with collections of different high-resolution satellite imagery such as Digital Globe's WorldView3 and Airbus Pleiades. The scenes may be level 2 products that have already been orthorectified or may exist as level 1 products that require orthorectification to a terrain model and possibly also pan sharpening to fuse the higher-resolution panchromatic imagery with the lower-resolution multispectral imagery. All such imagery can be added to a mosaic dataset with functions defining the processing to be applied on the fly. These functions can be chained together and include processes such as atmospheric correction, color balancing, and seamline-based blending.

Massive collections of aerial and drone imagery can be added to a mosaic dataset. Imagery from such aerial cameras could be preorthorectified but can also be the images directly from the frame cameras. As with satellite imagery, the parameters for orthorectification and pan sharpening of the imagery can be defined and the processing applied only if and when required. The ability to orthorectify and pan-sharpen the imagery on the fly enables significant savings in storage and processing as well as providing access to the original nonrectified pixel data. This is especially useful for oblique imagery, which can then be accessed either as orthorectified, where the images fit with the underlying map, or in image space, where the image remains as an undistorted oblique and features from the map may be projected into the image.

Mosaic datasets work with all forms of elevation data. Elevation data in the form of gridded Digital Elevation Models (DEM) and Digital Surface Models (DSM) can be added, as well as lidar (LAS or zLAS) point clouds or terrain datasets that are used to define constrained surfaces. From elevation-based mosaic datasets, a range of derived products such as slope, aspect, and hillshades can be generated on the fly. Using a mosaic dataset, global datasets such as Shuttle Radar Topography Mission (SRTM) can be combined with national and local datasets collected by photogrammetric or lidar methods. The user of the mosaic dataset can access the data as a single, virtually blended dataset at any resolution. Alternatively, a user can query the mosaic dataset based on attributes and filter the result so as to work with a single dataset or subset. An example implementation may be found in the ArcGIS Online World Elevation Service (esri.com/elevation) composed of over 44,000 datasets and over 2.2 TB of data.

The management of scientific, multidimensional rasters is also possible. Such rasters are often stored in netCDF, HDF, and GRIB formats. These can be added to a mosaic dataset with each slice of the raster becoming a separate uniquely identified record in the attribute table. It is then easy to use the database capability of the mosaic dataset to access any individual slice, obtain temporal profiles, or create data cubes for powerful visualization and analysis.

Mosaic datasets can also be created from the ArcGIS Full Motion Video (FMV) add-in. Instead of each video frame being a separate record, the FMV add-in enables the creation of a mosaic dataset from key frames that are based on a time interval. In this way, a virtual mosaic of the imagery can be created, providing a quick overview of different missions, as well as the ability to filter and find suitable sections of footage.

Automation of Workflows

When working with large image collections, it is important to automate the workflows to ensure they can be implemented routinely and precisely. Esri provides documented best practices and a set of workflow templates with sample data. These templates can be applied first with the sample data to test the process, then applied on your own data sources. The templates are provided as Python scripts so that they can be easily modified if required and used in automated environments. The use of these scripts further simplifies the process of creating mosaic datasets and enables the creation of mosaic datasets and services that can be updated on a regular interval without human intervention. For more information on these workflow templates, see esriurl.com/ImageManagement.

Providing Accessibility

The value of imagery is significantly increased if it is accessible. With ArcGIS, imagery can be accessed in multiple ways according to the requirement of the end user. Mosaic datasets can be directly accessed through desktop applications but also served through ArcGIS for Server as dynamic image services. Image services are the optimized way to access imagery from remote or networked systems. These image services can be accessed by a wide range of desktop, web, and mobile clients. The services are not static. Instead, the client applications interact with the servers, making imagery and data requests. The server reads the required data, performs the processing, and returns the image(s) for display or further analysis.

In addition to providing access to the pixel values, image services provide a searchable catalog of the data. Applications can perform queries against the catalog to find and return metadata about the imagery. The applications can then directly access the imagery products as image services. In some implementations, mosaic datasets reference reduced resolution or highly compressed versions of the products. This enables a fast catalog search and preview capability without the need to have the source data accessible to the servers.

An example of such an implementation is the landsatlook.usgs.gov/ site provide by the United States Geological Service (USGS). This site provides fast search and access to the complete archive of over 4 million Landsat scenes. Users can zoom to any location, perform a search for Landsat scenes corresponding to specific criteria, and then immediately see full-resolution versions of the scenes as well as apply some image enhancements. If needed, users can select full scenes that can be downloaded using traditional FTP. In this example, the mosaic dataset was created from an existing database catalog of metadata that also referenced the millions of JPEG compressed TIFF files that are full-resolution renderings of the original scenes (RGB only).

Image services also enable users to export or download sections of data. An *export* request causes the server to process and return imagery for a specified extent and pixel size. The format, projection, and processing of the imagery are defined by the client application, and the output may result in a number of input images being mosaicked together. A *download* request results in the server extracting and returning the original pixel values without sampling and with the original metadata. Downloading also allows the clipping of the data and format conversion if required. Such download and export requests are usually only done by applications that require their own copy of the data. In addition to returning the pixel values, the server can also return detailed metadata about the source imagery and the processing performed.

Scaling to the Cloud

Traditionally, organizations have stored imagery on direct access storage (DAS), enterprise network attached storage (NAS), or storage area networks (SAN). Such storage systems are typically optimized for fast, low latency and block-level access through file systems such as Network File System (NFS). These *block* storage systems are relatively expensive to scale, resulting in significant costs for large volumes of imagery. The advent of cloud computing has promoted the abundance of *object* storage. Examples include Amazon S3 or the Azure Blob store. They enable large numbers of computers to access data simultaneously. Such object storage is available from most cloud infrastructure providers but can also be implemented on-premises and in hybrid cloud environments. Object storage provides simple storage for massive petabyte data collections with significantly lower storage costs while providing very high reliability. Object storage does not, though, provide the same levels of access performance or latency as block storage and cannot be effectively accessed through a file system, so it is typically not suitable for image processing applications. ArcGIS includes technology that overcomes these limitations. To optimize performance, the imagery should be converted to a cloud-optimized format such as Meta Raster Format (MRF) as it is transferred to the cloud. By minimizing the number and size of requests, as well as using tile-based caching, fast access is achieved for imagery using object storage in the cloud. ArcGIS can also access tiled formats such as TIFF or JPEG2000 stored in object storage. These techniques enable organizations to manage their imagery holdings on inexpensive cloud storage and then provide fast access to the imagery as image services. Making use of the elasticity in a cloud infrastructure also enables the systems to efficiently handle the often varying loads.

An example of such a cloud-based image service utilizing object-based storage is the Landsat website referenced earlier (esri.com/landsatonaws). This site provides global access to over 100,000 full-resolution multispectral Landsat 8 and GLS scenes to the public through image services. The image services are updated daily with the new Landsat 8 scenes using a process that is fully automated. Access to all the multispectral bands with the various band combinations and indices are provided. All the processing is being applied on the fly when the client requests access. The source multispectral Landsat (LIT) scenes are stored on Amazon S3 object storage as tiled GeoTIFF files as part of the Amazon Public Datasets. Each band of each scene is a separate file. The definition of how to process and blend the files is defined in the mosaic dataset. The imagery referenced by the mosaic dataset is then served as image services using elastic Amazon Machine Images (AMIs) that can scale up or down depending on the server load. The data is accessed and processed on demand, with the ability for the client applications to refine the processing they'd like to have applied. These image services can be used in a range of web and desktop applications to quickly visualize and analyze the Landsat dataset made available by USGS.

Optimal Data Transfer

To maintain the full information content, lossy compression is not typically applied on source data. However, there are many advantages to using compression while transmitting server-side processing results to client applications, as it considerably reduces the volume transmitted. Lossy compression, such as JPEG, is typically used for continuous imagery data, while lossless PNG is used for discrete data (e.g., a categorical dataset generated through image classification). These formats are used as they can be displayed directly in web browsers. Although multiple MB of data may need to be read and processed by the server, typically only about 200 KB of data is returned for each user's pan/zoom operation. Using compression ensures fast access. The client applications also have the ability to change the compression type and quality used to

transmit the results. For example, users may want a lower quality (higher compression) while they navigate and a higher quality once they've found their area of interest and want to perform detailed interpretation. This client-side changeable compression for transmission enables the use of image services over very poor bandwidth connection, yet users can still access the full information content to perform analyses.

Some applications need to provide highly dynamic user experiences, such as the creation of interactive height profiles or viewsheds from elevation. Such applications cannot rely on the server to perform the required rendering. Instead, they need access to the full pixel values for an area of interest and then use client-side processing to perform the rendering. ArcGIS image services can return such data values to the applications. To reduce the data transfer volumes, LERC compression is used, as it can handle any bit depth and defined error tolerances. The LERC compressed data can be read and processed by JavaScript- and HTML5-based applications, enabling the use of imagery in modern web applications, as well as in distributed big data processing systems.

Open Standards

ArcGIS for Server supports Open Geospatial Consortium, Inc. (OGC), Web Map Service (WMS), Web Map Tile Service (WMTS), Web Coverage Service (WCS), Web Processing Service (WPS), and KML standards. These OGC standards can be used to access imagery in all OGC-conforming applications to aid in interoperability. Applications based on ArcGIS, such as ArcGIS for Desktop, can also consume OGC services. For users who want additional features beyond the simple display, data access, and temporal control provided by OGC standards, Esri provides the ArcGIS REST API. Most of the more advanced client applications use the REST API, which is also easily integrated into applications.

Client-Defined Processing

Image services enable advanced applications that run on lightweight browsers and mobile devices to display imagery and raster data while fully utilizing the server-based processing. Client applications use the image services API to define how imagery is accessed and processed by the server. Users within the applications can pan and zoom to any extent and scale while the server processes and returns the imagery based on the different renderers or analytic functions. The client can also specify the coordinate system desired, and the server will reproject on the fly directly from the source coordinates to the output without the need for temporary intermediate files.

Each image service can have a list of functions associated with it, and applications can select the function to apply as well as parameters of the function. For example, a user can select a hillshade function on elevation data and define a custom azimuth and sun angle to use. If allowed by the server administrator, the client applications can also transmit function chains to the server. The function chains can define collections of functions to be processed by the server. If the administrator of the service has published Python-based raster functions with service, then the client application can transmit parameters to be used by these functions. The client applications can therefore specify and control a wide range of processing to be applied on the source rasters.

The image service APIs are very rich and enable a wide range of processing parameters to be set. For example, the sampling method to be used can be set to either nearest neighbor, bilinear, or cubic convolution. The API is also able to make requests that return data computed from imagery. For example, if the source imagery has sensor orientation information, the client can perform mensuration to determine the heights of objects. The

application can request a user to define the displayed location of the top and base of a building, then transmit these to the server to compute and return the building height. Similarly, image services can perform tasks such as computing and applying signatures for supervised classification or returning statistics and histogram information for polygons such as field boundaries. These client-defined processing functions enable the development of applications that invoke advanced image processing capabilities beyond just visualizing the data.

Creating and Serving Tile Caches

Some applications only require imagery as a simple background for visual context. This can be easily achieved via dynamic image services, but in cases where such access needs to be scaled to large numbers of users, it is often advantageous to serve imagery as a static tile cache, which minimizes the server overhead. Most consumer web applications that provide background imagery use the static tile cache method to serve background imagery. In workflows to create such tile caches, mosaic datasets are used to manage the data sources, define processing to be performed, and perform QC of the expected output. The tile cache can then be generated either in ArcGIS for Desktop or, for large projects, using ArcGIS for Server to run the process in parallel using multiple machines. This generates cached tiles that enable the server to quickly return the imagery when it is requested with no additional processing. Applications obtain from the server the schema of the tile layout and can request the tiles required to cover the screen at the appropriate scale. In addition to the option for preprocessing all tiles, ArcGIS for Server provides the option to generate tile cache on-demand, so it is only created when an area is first viewed by a user.

Geoprocessing

There are many applications that need to access and process non-raster-based geospatial data, such as vector-based road networks, or access and process raster data that exists beyond the local geographic extent of a request. Examples include the computation of downstream water trace or watersheds from elevation models. For such processes, ArcGIS geoprocessing tasks can be served through ArcGIS for Server. Geoprocessing services have access to the very extensive range of geoprocessing tools within ArcGIS as well as tools developed by Esri partners. Unlike image services, geoprocessing tasks are asynchronous and can return all forms of geospatial data (e.g., imagery, tables, or maps). A task may return a result very quickly, but for larger processing tasks, the response may take a longer time and so the asynchronous requests are more appropriate. Such geoprocessing tasks can access raster data directly as files, from a mosaic dataset, or from an image service. Geoprocessing can be used to perform simple tasks such as the computation of a watershed but are also used as the interface to big data analytics, where a geoprocessing task can initiate a set of processing tasks on many machines, then compile and return the results.

Raster Analytics and Big Raster Processing

Mosaic datasets, along with image services, can be used to provide access to massive collections of imagery and raster data. They are significantly changing the way imagery is managed and processed by enabling the source data to be stored once with the ability to quickly create multiple imagery products generated on demand. The processing performance is typically much faster than the traditional approach of passing imagery through multiple processing steps with multiple inputs and outputs. In ArcGIS, the processes are concatenated together through very efficient algorithms that eliminate the creation of intermediate datasets, which were throughput bottlenecks that consumed disk space. As a side benefit, the imagery quality is improved because of the reduction in the

number of times the data is sampled. Using ArcGIS, only a single sampling of the imagery occurs between the source and the data displayed on the screen.

There are processes, such as the generation of tile caches, tiled orthophotos, or performing detailed image classification on large areas, where the ArcGIS servers are used to preprocess the data. These processes are scalable and can be performed quickly as the tasks can be distributed over multiple machines. By combining the raster and vector processing capabilities of ArcGIS, these tasks can also be applied to a wide range of raster and spatial analysis tasks.

Conclusion

ArcGIS is a comprehensive enterprise platform for image management, processing, analysis, and serving. With ArcGIS, users can create mosaic datasets that catalog the imagery and rasters within an organization while defining the processing to be applied. This makes imagery accessible to anyone who needs it in a timely manner. These mosaic datasets can be used directly with ArcGIS for Desktop or ArcGIS Pro or served as image services to a wide range of web and mobile applications through ArcGIS for Server. Those applications can access the imagery as useful information products, enabling the full information content of the imagery to be exploited.

Getting Started

Esri has many ways to learn more about managing imagery with ArcGIS:

- Extensive online help documentation—esriurl.com/WhatIsRasterData
- The ArcGIS Imagery resource center—esriurl.com/ImageryResourceCenter
- Image Management Workflow web pages with links to best practices and automation templates for working with different imagery—esriurl.com/ImageManagement

Ready to get started or have additional questions? Email the imagery sales and technical team directly at Imageryinfo@esri.com.

The Esri Professional Services team is also available to jump-start projects and help organizations quickly get started. For more information, check out esri.com/services/professional-services/rent-a-tech.



Esri inspires and enables people to positively impact their future through a deeper, geographic understanding of the changing world around them.

Governments, industry leaders, academics, and nongovernmental organizations trust us to connect them with the analytic knowledge they need to make the critical decisions that shape the planet. For more than 40 years, Esri has cultivated collaborative relationships with partners who share our commitment to solving earth's most pressing challenges with geographic expertise and rational resolve. Today, we believe that geography is at the heart of a more resilient and sustainable future. Creating responsible products and solutions drives our passion for improving quality of life everywhere.



Contact Esri

380 New York Street
Redlands, California 92373-8100 USA

1 800 447 9778
T 909 793 2853
F 909 793 5953
info@esri.com
esri.com

Offices worldwide
esri.com/locations