

IMAGERY SUPERPOWERS

Raster analytics expands imagery use in GIS

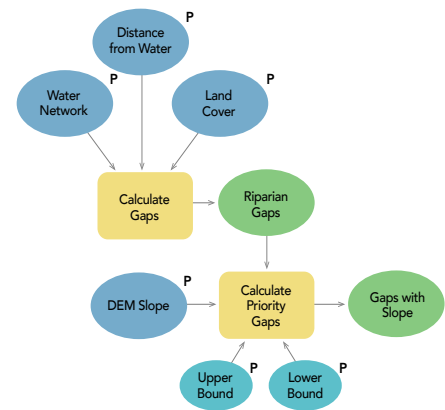
By Jeff Liedtke, Product Engineer, Esri Raster Team

Aerial, drone, and satellite imagery is more available than ever before and can add valuable information and context to GIS projects. However—before the raster analytics capabilities in the ArcGIS platform—incorporating massive imagery workflows into your GIS could be overwhelming.

Esri developed raster analytics capabilities so that you can quickly process large imagery collections and extract and share meaningful information for critical decision support. The capabilities of ArcGIS Enterprise are extended with ArcGIS Image Server configured for the roles of raster analytics and dynamic image services to host distributed processing and storage. In this scalable environment, you can implement computationally intensive image processing that used to be impractical, impossible, or cost-prohibitive.

Raster analytics can be run locally to fully utilize your existing ArcGIS Image Server on-site. However, using the scaling

capabilities of Raster Analysis Server and Raster Image Server in a distributed cloud computing environment maximizes efficiency by exploiting the elastic processing and storage capacity of cloud computing and storage platforms such as Amazon Web Services (AWS) and Microsoft Azure. With this elasticity, you can increase or reduce capacity depending on the size and urgency of projects. Image processing and analysis jobs that used to take days or weeks can be completed in minutes or hours, bringing imagery projects that were impossibly large or daunting within reach. (See the accompanying article, “Tackling a Monumental Project.”)



↑ Raster function chains can be interactively created in the ArcGIS Pro Raster Function Editor pane.




With this system, you can apply standard raster processing tools and functions in ArcGIS, build custom functions and tools, and combine multiple tools and functions into raster processing chains to execute custom algorithms on large collections of raster data. Results are stored, automatically published, and shared across your enterprise.

This flexibility in processing and storage capacity lets you meet demanding project timelines and tight cost requirements. Resources can be ramped up to meet a surge in demand, then released when no longer needed.

Power, Extensibility, and Availability

To enable raster analytics, ArcGIS Image Server is configured to perform distributed computing in a processing and storage environment that maximizes processing speed and efficiency. Built-in tools and functions cover preprocessing, orthorectification and mosaicking, remote-sensing analysis, and an extensive range of math and trigonometry operators. These functions process multispectral and hyperspectral imagery, multidimensional scientific data addressing data cube type problems, and temporal data for time series analysis.

Advanced raster analysis processing chains can be built from the more than 200 functions and tools that are available



↑ The raster analytics processing workflow

out of the box and saved as raster function templates.

You can write custom functions that extend the platform's analytical capabilities and deploy custom image processing algorithms to address unique applications and problem sets. Developers can build production processing workflows by utilizing the ArcGIS Enterprise REST API, ArcGIS Server REST API, or ArcGIS API for Python with raster function objects using code samples found on the ArcGIS for Developers (developers.arcgis.com) website or scripts available from the Esri repository under the `arcgis.raster.functions` module on GitHub.

Custom processing chains can also be


built using the Python Raster Function that allows you to write image processing algorithms in Python and execute them in a distributed environment. (See github.com/Esri/raster-functions/wiki/PythonRasterFunction.)

Raster analytics integrates image processing and analysis with the rest of the ArcGIS platform, streamlining and simplifying collaboration and sharing. Users across your enterprise can contribute data, processing models, and expertise to an imagery project, and the results can be shared with individuals, departments, and organizations in your enterprise.

Removing Barriers

Efficient and powerful processing of large image and raster data collections using compute-intensive processing chains allows users to address projects and jobs that were out of reach in the past. This lowers the bar for smaller or midsize companies, allowing them to undertake projects that were exclusively the province of large, well-resourced organizations. Limited processing and storage resources are no longer barriers for organizations that want to take advantage of opportunities that have demanding image processing requirements and delivery time frames.

Elasticity in the deployment of raster analytics provides scalability in terms of speed as well as dataset size. Sometimes the requirement is that data be quickly processed for applications such as emergency management support, processing multidimensional scientific data from sensors, or constant monitoring by drones.



Another important benefit of this ability to scale rapidly is surge capacity. This has always been an operational challenge for contractors and service providers. Resources and costs can now be tailored to specific projects. Leveraging cloud resources enables an organization to address acute capacity needs without high capital investment, maintenance costs, or risks.

When an organization wins a project, it can quickly ramp up to a level appropriate to the project's time frame and budget, complete the project, and then release resources when they are no longer needed. Experience estimating time and cost for certain types of projects will become more refined over time, making an organization more competitive and lowering risk and liability.

Raster Analytics under the Hood

Image Server configured for the role of raster analytics provides software and user interfaces to organize and manage processing, storage, and sharing of image and raster data, maps, and other geographic information on a variety of devices. This integrated system manages the dissemination and storage of results on-premises and behind the firewall for classified deployments, in cloud processing and storage environments,

or in a hybrid environment that uses a combination of on-premises and cloud.

The foundation of raster analytics is ArcGIS Enterprise, which includes an enterprise GIS portal, ArcGIS Data Store, and ArcGIS Server, as well as Image Server configured for raster analytics, raster data store, and ArcGIS Web Adaptor. ArcGIS Enterprise integrates the components of the raster analytics system to support scalable, real-world workflows.

Scale powerful processing and storage capabilities by deploying ArcGIS Enterprise in the cloud via AWS or Microsoft Azure. For example, you can automatically scale capacity up and down according to conditions you define or automatically disperse application traffic across multiple instances for better performance. Esri makes deployment easier by providing Cloud Builder for Microsoft Azure or AWS CloudFormation with sample templates to configure and deploy your system in the cloud.

Running a Raster Analysis Job

Using the user interface in ArcGIS Pro or web map viewer, running raster analysis jobs is straightforward. You simply submit your job, and the way you defined and set up your portal determines how the processing

and storage is distributed across the available servers, instances, and data stores.

The raster analysis workflow begins with preparing your image and raster data for efficient processing. This involves registering your data and converting it to a distributed version for optimized, distributed processing and storage, but distributed storage is not mandatory. The data is converted into Cloud Raster Format (CRF), a tiled format that is optimized for performance.

Raster analysis will work with your existing data, but it will not be as efficient as when you use optimized data. Optimal performance is achieved when your source data has been prepared and staged in the data store. Tools are provided to convert your data files into a distributed version of the data, which optimizes it for efficient distributed processing and the storage of published results.

Develop, test, and optimize raster processing chains using the more than 200 functions and tools in ArcGIS Pro or web map viewer. With the raster function editor in ArcGIS Pro, you can interactively build and test a processing chain, preview results, and verify those results in a dynamic processing environment. With ArcGIS Pro, you can also create customized functions using Python.

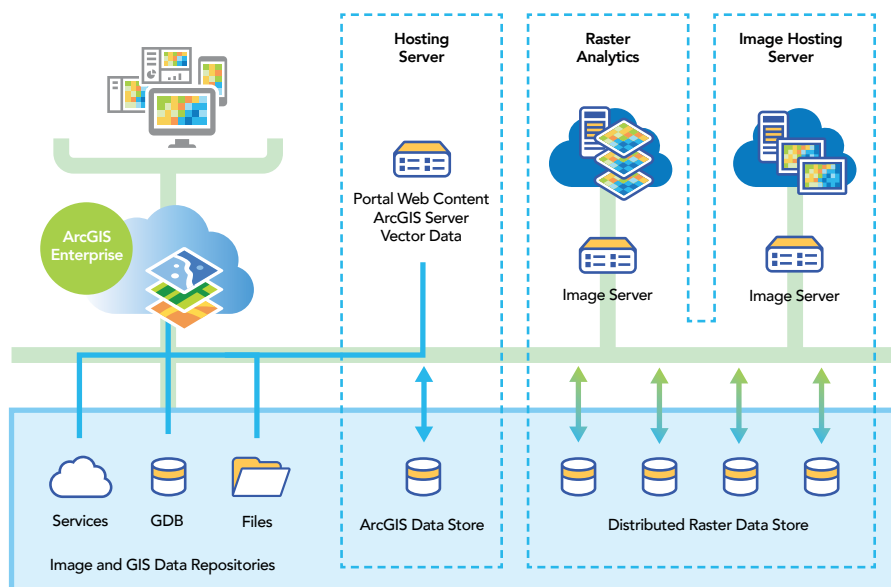
When you are satisfied with your processing chain, you can save it as a raster processing template. Processing templates can be shared with your organization through ArcGIS Enterprise. Once a processing chain has been optimized, submit it to the ArcGIS portal, which manages the distribution of processing, storage, and automatic publication of results.

You submit a raster analysis processing job by specifying the portal and naming the output web image layers. The job is sent to the Raster Analysis Server for distributed processing using your processing chain and the raster data prepared and stored in the raster data store repositories.

The raster output from the Raster Analysis Server is *automatically* published to the distributed Raster Data Store, while any vector output is published to the GIS Data Store. Since the analytical results from your raster analysis job are automatically published, they are immediately available to your enterprise for sharing and further analysis.

Raster analytics components and relationships

ArcGIS Server with Image Server is the combination of licenses that enables creating and hosting of image services (Image Hosting Server). When an ArcGIS Server with Image Server is federated with an ArcGIS Enterprise server (Hosting Server), the raster analytics capabilities become available. These servers use ArcGIS Server data stores to store and access the results of raster analytics output.



Deploying Image Server

The ideal deployment of Image Server with raster analytics is composed of three (or more) server sites that perform the primary roles of the portal host server, raster analysis server, and image hosting server. Two licenses are required for raster analytics: ArcGIS Enterprise and Image Server.


The hosting server is your portal's server for standard portal administration and operations such as managing and dispensing processing, storage, and publication of results to raster analysis servers, image servers, and data stores. It also hosts ArcGIS Data Store for GIS data and allows users to publish data and maps to a wider audience as web services.

Raster analysis jobs are processed by image servers dedicated for raster analytics, composed of one or more servers, each with multiple processing cores. The image processing and raster analytics tasks are distributed at the scene or subscene level depending on the tools and functions used. Image Server manages the processing results to either ArcGIS Data Store on the hosting server for feature data products or to the raster data store for imagery and raster data products. The raster data store can be implemented using distributed file share storage or using cloud storage such as Amazon S3 or Microsoft Azure Blob Storage.

The image hosting server hosts all the image services generated by the raster analysis server. It includes the raster data store and returns results requested by members of your enterprise.

Templates and system configuration apps assign the roles of the servers and data stores and also set the permission structure for all the users across your enterprise. This facilitates optimal flexibility in configuring and implementing your raster analytics system to address specific projects. Multiple servers can be scaled up for raster analytics processing and storage as required.

This architecture can be varied according to your requirements and budget. For instance, the image hosting server and the hosting server may be the same server. The image repositories can be local drives on the ArcGIS servers, already in one of the server's data stores, or in a cloud repository such as an Amazon S3 bucket where Amazon AWS hosts repositories of



published imagery and raster data. (For the latest information on how to configure ArcGIS Enterprise for raster analytics, refer to the online help for ArcGIS Enterprise.)

Making the Previously Impossible, Possible

Because raster analytics is integrated with and takes advantage of, components of the ArcGIS platform, it utilizes the rich set of analytical capabilities of ArcGIS Desktop and turbocharges processing by distributing the storage and processing of rasters. These capabilities create a powerful, distributed raster processing, storage, and sharing system that can be scaled to fit any project having demanding requirements.

A paradigm shift in imagery use has occurred. With raster analytics, organizations—both large and small—are empowered to address projects with diverse and demanding requirements that were previously impractical by scaling operations in an efficient and cost-effective manner.

MapLogic
Map book solution for ArcGIS®

Take your map book creation to the next level!

- A complete map book in a single MXD
- Multiple layouts in ArcMap
- Automated street indexing



Download a fully unlocked evaluation version and try it out for 30 days

www.maplogic.com

407-657-1250 • sales@maplogic.com