ESRI
EXTERNAL

ArcGIS Image Dedicated User Manual

July 23, 2026

380 New York Street
Redlands, California 92373-8100 USA
909 793 2853
info@esri.com
esri.com

esri | THE SCIENCE OF WHERE™

Copyright © 2026 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, The Science of Where, 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.

Contents

What is ArcGIS Image Dedicated? 6

Setup 6

Registration 6

Cloud Storage 7

Data Access 7

ACS Connection 7

Raster Proxies 8

Set up Tiled Imagery 9

Set up Dynamic Imagery 9

Imagery Server Performance Metrics 9

Upgrade Dynamic or Tiled Imagery 10

Set up an RDBMS 11

Enable Processing and Analysis 12

Control Access to Cloud Storage 12

Grant Secure Access to Data in AWS 12

Grant Secure Access to Data in Azure 14

Private Storage Connections 14

Concepts 15

ArcGIS Image Dedicated Website Overview 15

What is an Image Service? 15

Set Up 15

Dynamic Imagery 16

Overview 16

Tiled Imagery 16

Overview 16

Manage Cloud Stores in an Imagery Server 16

Data Storage Account 16

Cloud Storage Outside the Data Storage Account 17

Cloud Store Management Actions 18

Service Management Toolbox 19

Create Image Definition Tool 20

Overview 20

Parameters 20
Get Job Status 21
Overview 21
Parameters 22
Manage Dynamic Imagery 22
Overview 22
Parameters 22
Manage Tiled Imagery 25
Overview 25
Parameters 25
Set System Service Settings 27
Overview 27
Parameters 27
Processing and Analysis 28
Overview 28
Enabling or Disabling a Processing and Analysis Service Task 28
Output to Cloud Storage 29
Publish Results 29
Managed RDBMS 29
Overview 29
Create an RDBMS 30
Access an RDBMS 30
Manage RDBMS Access 30
Use an RDBMS 31
Extend the Term Length of a Running RDBMS 31
Upgrade an RDBMS 32
Pro Machines 33
Overview 33
Create a Pro Machine 33
Access a Pro Machine 36
Using a Pro Machine 37
Stop a Pro Machine 37
Pro Timer 38

Reset Password 38
Delete a Pro Machine 38
Administration 39
Email Notifications 39
Usage Reports 39
Control Access Items 39
Admin Access Item 40
Server Access Item 40
P&A Access Item 40
Designate a Group for ArcGIS Image Dedicated Access 41

What is ArcGIS Image Dedicated?

ArcGIS Image Dedicated is a managed Software as a Service (SaaS) offering from Esri, designed for organizations that require imagery to be stored and managed within their own cloud environments or sourced from public datastores. It provides a scalable and secure platform for advanced imagery management, processing, and analysis.

Key Capabilities

With ArcGIS Image Dedicated, you can:

- Publish tiled imagery for fast access using formats like Cloud Raster Format (CRF), tile cache, raster tiles, vector tiles, 3D tiles or scene layers.
- Publish dynamic imagery with server-side mosaicking and real-time processing.
- Run processing and analysis services to convert data or extract insights from imagery.
- Run Pro Machines, virtual desktops preloaded with ArcGIS Pro .
- Create a managed RDBMS, allowing you to author mosaic datasets within an enterprise geodatabase and publish dynamic imagery.

Supported Cloud Platforms

ArcGIS Image Dedicated supports imagery stored in the following cloud environments:

- **Amazon Web Services (AWS)**
- **Microsoft Azure**

Setup

Registration

To use **ArcGIS Image Dedicated**, an active **ArcGIS Online organization** is required. To register your organization, visit ArcGIS Image Dedicated site and follow the registration instructions.

Note: You need Administrator privileges to register the organization.

If you're logging in for the first time, you will be redirected to the user registration page. The point of contact and email information will be automatically populated based on your ArcGIS Online organizational account details.

You are prompted for ArcGIS Image Dedicated permission to create an admin access item in your ArcGIS Online account. The admin access item is used to administer privileges to your organization's ArcGIS Image Dedicated subscriptions. The ArcGIS Image Dedicated Admin item will be available as an application within the list of items on the Content tab of ArcGIS Online. The following image depicts the same.

Once registered, the administrator can share access of the admin item to other members in the organization. These members are now administrators to all the ArcGIS Image Dedicated services.

Note: Access to admin item should be reserved and assigned only to those organization members who need full administrator access. Create other items to grant access to finer-grained services.

Cloud Storage

With ArcGIS Image Dedicated, you can leverage raster analysis and stream directly from your imagery housed in cloud storage accounts. Currently, AWS S3 buckets and Microsoft Azure containers are supported as cloud repositories for imagery and raster data. The imagery and raster data can either be used directly as datasets to publish image services or they can be added to a mosaic dataset, which can also be used to publish dynamic image services and as the source for analysis.

If your imagery data does not currently reside in cloud storage, depending on your cloud storage service, you can transfer a copy of your local data to the cloud. The Transfer Files geoprocessing tool in ArcGIS is designed to streamline the movement of data between your local environment and cloud storage using an ArcGIS Cloud Storage (ACS) file. There are other alternatives for transferring data to the cloud, including third-party tools and tools like OptimizeRasters. For transferring extremely large amounts of data, we suggest utilizing services offered by cloud providers like Amazon Snowball. Moreover, for data transfer, AWS CLI is a robust command-line utility that enables interaction with multiple AWS services, including the efficient transfer of files to Amazon S3. For individuals utilizing Azure storage accounts, Microsoft Storage Explorer offers an easy-to-use graphical interface for file management and transfer, serving as an intuitive substitute for command-line tasks.

Data Access

Rasters in cloud storage can be accessed using the following methods:

- Using an ArcGIS Cloud Storage (ACS) connection to access data
- Using raster proxies to access data

ACS Connection

An ArcGIS Cloud Storage (ACS) file is a connection file for cloud storage that supports ArcGIS. This file allows users to connect directly to cloud storage, providing access to raster data stored in the hosted cloud. The ArcGIS Cloud Storage (ACS) file streamlines data access and enhances workflow efficiency by:

- **Local-like access**: Browse and interact with cloud-hosted raster datasets as if they were stored on your local machine.
- **Flexible security**: Supports multiple cloud security profiles, allowing access to diverse imagery collections from a single environment.
- **Encrypted authentication**: Credentials are securely encrypted within the ACS file, ensuring safe and reliable access.

- **Optimized performance**: Enabling the Pixel Cache option on the Add Rasters dialog when authoring mosaic datasets accelerates repeated data retrieval by minimizing cloud storage checks.
  Note: When the mosaic dataset is to be used only for analysis, disabling pixel caching is recommended to avoid unnecessary overhead.

We recommend ACS connection file when working with dynamic datasets that frequently change and require instant access.

How to use?

- Create the ACS file using the Create Cloud Storage Connection File geoprocessing tool.
- Access the Catalog pane in ArcGIS Pro, resembling a local file system, to browse and choose files for inclusion

Note: Performance will be significantly degraded if the cloud storage and the ArcGIS Image Dedicated server are in different regions. To avoid delays in accessing imagery and processing requests, ensure your server and storage are in the same region.

Support of private storage connections in ArcGIS Image Dedicated allows ACS Files using AWS cloud provider to be created without access credentials. Additionally, the support for Azure Front Door Endpoint is added in ArcGIS Image Dedicated. The endpoint can be used in creating ACS Files and Data Storage Accounts.

To read more about it refer to the Data Storage Accounts and Private Storage Connections Guide and ArcGIS Image Dedicated Service Management Geoprocessing Toolbox documentation available under Resources in the ArcGIS Image Dedicated site.

Raster Proxies

Raster proxies are small XML files stored on local file systems that serve as a virtual pointer to larger raster datasets stored in cloud storage. They are used to locate and access the actual data from the cloud as needed. This approach allows ArcGIS to optimize performance and enhance data handling in several ways:

- Uses a small XML file to reference cloud-based raster data instead of downloading large raster files
- Enables efficient data transfer into and out of cloud storage
- Retrieves only the required raster tiles from the cloud, reducing storage and egress costs
- Supports most GDAL-readable raster formats and any raster file extension
- Can be used directly in ArcGIS or added to a mosaic dataset for flexible integration
- Improves display and processing speed by caching previously accessed pixel data and tile locations locally

Recommended when working with massive cloud-hosted raster datasets that exceed local storage capacity.

How to use?

- Use the OptimizeRasters tool to create Raster Proxy files, select cloud-stored raster data, and save the proxy files locally.
- Add these Raster Proxies to your map or mosaic dataset for faster display and processing by caching pixel data locally.

Set up Tiled Imagery

As part of the tiled imagery subscription, you can use dedicated servers to serve tiled imagery services. Follow these steps to request and add a new server:

1. Navigate to the homepage of the ArcGIS Image Dedicated portal.
2. Access the Server Request Form
    - Select the Tiled Imagery tile.
    - Click on **+ New Server** to open the server request form.
3. Fill Out Server Details
    - Specify the cloud provider (e.g., AWS or Azure) where your imagery is stored.
    - Choose the cloud region for server deployment.
        - This region should match the region of your cloud storage to ensure optimal performance.
4. Submit the Form
    - Review your entries and submit the form.
5. Server Creation and Confirmation
    - Your server will be created based on the specifications you provided.
    - Once active, it will appear in the Active Servers list.

Set up Dynamic Imagery

As part of the dynamic imagery subscription, you can use dedicated servers to serve dynamic image services.

1. Navigate to the homepage of the ArcGIS Image Dedicated portal.
2. Access the Server Request Form
    - Select the Dynamic Imagery tile.
    - Click on **+ New Server** to open the server request form.
3. Fill Out Server Details
    - Specify the cloud provider (e.g., AWS or Azure) where your imagery is stored.
    - Choose the cloud region for server deployment.
        - This region should match the region of your cloud storage to ensure optimal performance.
4. Submit the Form
    - Review your entries and submit the form.
5. Server Creation and Confirmation
    - Your server will be created based on the specifications you provided.
    - Once active, it will appear in the Active Servers list.

Imagery Server Performance Metrics

Imagery Server Performance Metrics (available in Imagery Server version 11.5 onwards) provides real-time insights into the resource usage of each machine powering your imagery server within an Auto Scaling Group. It offers detailed, per-instance graphs for average CPU, Memory, and Disk utilization,

updated every minute through a lightweight monitoring agent. Server-level metrics such as the total number of requests served by the imagery server, give a complete view of workload and service demand. This feature is designed to help you monitor imagery server workloads, optimize performance, and maintain seamless service delivery. You can benefit from:

- Real-Time Monitoring: Track per-instance CPU, memory, and disk utilization alongside overall request volume to understand usage intensity and server demand.
- Customizable Analysis: Filter metrics by resource type, instance, and time range, and export data to CSV for reporting, auditing, or planning.
- Historical Insights: Access up to 7 days of past metrics, including three recently scaled-in instances, to evaluate resource consumption and refine scaling decisions.
- Organized Metrics for Clarity: Metrics are grouped by instance and load balancer, making it easy to interpret performance data and correlate server demand with resource usage.
- Performance Optimization: Identify peak usage periods and make data-driven scaling and resource allocation decisions.

To access the metrics, click on the … for the server whose performance you want to view, then select View Performance Metrics from the drop-down list. This opens the dashboard where per-instance CPU, memory, and disk usage, along with server-level request volume, can be monitored in real-time and analyzed over historical periods.

Upgrade Dynamic or Tiled Imagery

As part of the dynamic or tiled imagery subscription, you can now request an upgrade of your existing server to the latest version.

1. Access the Server Upgrade Request Form

    - Navigate to the Active Servers list.
    - Click on … for the server you wish to upgrade.
    - Select the Upgrade Server Version.

2. Complete the Upgrade Form.

    - Server Alias (auto populated)
    - DNS Prefix for the new server (auto populated)
    - Imagery Server Version you wish to upgrade to (enter the required details)

Submit the form.

3. Server Upgrade Processing

    - Your request will be processed, and the upgraded server will be created.
    - Details of the new server will be sent to you via email.

4. Trial Period and Verification

- The new server will remain active for a 7-day trial period.
- During this time, you can:
  - Test and verify services.
  - Respond via email to confirm the DNS swap or request a hold if issues are found

5. Next Steps

- If no response is received within 7 days, the DNS swap will proceed automatically.
- After the swap, the old server will be deleted.

Set up an RDBMS

As part of the Processing and Analysis subscription, you can now request an RDBMS to author mosaic datasets in an enterprise geodatabase and use it to publish dynamic imagery. Check the requirements and details.

To create a new RDBMS

1. Navigate to ArcGIS Image Dedicated website

- Open the Processing and Analysis page.

2. Access the RDBMS Section

- Click on the RDBMS tile.
- Select + New RDBMS to begin the request.

3. Complete the RDBMS Request Form

Fill in the required fields:

- Cloud Provider (e.g., AWS, Azure)
- Deployment Cloud Region
- Must match the region of your dynamic imagery server
- RDBMS Term Length (e.g., 3 months, 6 months, 12 months)
- RDBMS Type (e.g., PostgreSQL)
- Alias (custom name for easy identification)
- Database Name

4: Submit the Request

- Review all input details.
- Click Submit to initiate RDBMS creation.

5: RDBMS Provisioning

- The system provisions the RDBMS based on your specifications.

6: Verify Deployment

- Once created, the RDBMS will appear in your Active RDBMS List.

Enable Processing and Analysis

As part of the Processing and Analysis subscription, you can enable tasks that allow you to process new imagery to create new representations of the data or perform analysis to extract information from imagery. Such tasks include format conversion, automated creation of mosaic datasets, classification, and deep learning.

Control Access to Cloud Storage

ArcGIS Image Dedicated allows you to create image services on Esri-managed infrastructure, with the data stored in cloud storage. To ensure Image Dedicated servers have access to read or write your data, permission needs to be set in the cloud storage infrastructure. Currently, ArcGIS Image Dedicated supports only two cloud regions: AWS and Azure.

Grant Secure Access to Data in AWS

When creating an ArcGIS Image Dedicated server, you can specify a data storage account as part of the setup. Doing so grants the server automatic read and write access to the designated cloud storage—whether it's Azure Blob Storage or Amazon S3—without requiring manual configuration.

This direct access eliminates the need to copy ACS files into the server environment, streamlining your workflow and reducing data redundancy. All requests to the cloud storage are authenticated through the data storage account, and manual ARN setup is not required since access is granted directly through the provided storage credentials.

When using external buckets without a data storage account, you must set up an ARN to grant access. Users must also copy the ACS file to enable authentication when accessing the bucket. If the ACS file is not copied, any published services using the external bucket will fail, and imagery may appear as a black-and-white pattern due to missing permissions.

To set up an Amazon Resource Name (ARN) for ArcGIS Image Dedicated:

1: Download Setup Resources

- Go to the **ArcGIS Image Dedicated website**.
- Locate the **Downloads tile**.
- Download the **Create Storage Account** zipped folder.
    - This folder contains both **Python scripts** and **AWS CloudFormation templates**.

2: Choose Your Setup Method

- You can use either:

- **Python scripting** for manual control and customization.
- **AWS CloudFormation templates** for automated deployment.

3: Create Secret ARNs

- Use the provided script or template to generate **Amazon Resource Names (ARNs)**.
- These ARNs grant secure access to your **S3 buckets** for:
  - Reading imagery
  - Writing analysis output
  - Storing ArcGIS server logs

4: Configure Access Permissions

- The ARNs are configured to grant access **only to Esri-owned AWS accounts**.
- You retain **full control** over:
  - When to **revoke access**
  - How to **modify permissions**

5: Validate and Monitor

- Confirm that the ARNs are correctly applied and functioning.

The Python script automates the process by creating a secret in AWS Secrets Manager for the given Identity and Access Management (IAM) user keys. It then encrypts them using Amazon's Key Management Service (KMS) and adds the policy to the resource to allow the Esri account to access the key. An ARN will be generated and shared only with the machines allowed to perform the processing. The ArcGIS Image Dedicated servers can retrieve the Key with the generated ARN.

Note: To use the Python script, you will need Python 3 with boto3 installed and AWS CLI configured with keys.

Alternatively, you can upload the CloudFormation template to AWS Console to create a secret ARN. A CloudFormation template is used in AWS to allocate services in a template format. The templates must be run in the same region as the ArcGIS Image Dedicated server. If you already know which S3 buckets you want to grant access to, use the aid-iam-ssm.template to configure permissions. This template will create an IAM user and allow access to the specified bucket and store the keys in AWS secret manager encrypted by KMS. If you only want to create the secret manager and KMS, you can use the aid-ssm.template, which stores the keys provided in the AWS secret manager encrypted by KMS.

After you've successfully implemented the ARN with either the script or the CloudFormation template, the "Secret ID" will contain the ARN of the shared secret, which can then be used to create the data storage accounts using ArcGIS Image Dedicated.

Grant Secure Access to Data in Azure

If your data is housed in the Azure cloud infrastructure, you need to be able to grant secure access to your data. In Azure, there are two ways to create storage accounts—Access keys and Shared Access Signature (SAS) tokens.

Access keys are used to authorize access to data in a storage account through a Shared Key authorization. Creating a data storage account using access keys grants ArcGIS Image Dedicated servers full permissions—this includes, but is not limited to, the ability to create, read, update, and delete files, queues, blobs, and tables.

Alternatively, an SAS is a URL that grants restricted access rights to Azure storage resources. By distributing an SAS URL to these clients, you grant access to a resource for a specified period. To restrict granted access, you can generate an account-level SAS token; the permissions can be customized according to the client.

After you've successfully submitted a data storage account request, the Azure storage Account key or SAS is encrypted using Advanced Encryption Standard 256, which can then be used by ArcGIS Image Dedicated.

Private Storage Connections

The storage access methods described in the previous sections use storage credentials (such as Shared Secret ARN, Azure Account Keys, or SAS Tokens) to securely access user-managed cloud storage through the cloud provider's public storage endpoints. Although these connections are authenticated and encrypted, the storage endpoints remain publicly routable.

For users who require network isolation and private connectivity, ArcGIS Image Dedicated provides Private Storage Connections. Private Storage Connections enable secure communication between Image Dedicated resources and user-managed cloud storage over the cloud provider's private network, eliminating the need for data traffic to traverse the public internet.

Private Storage Connections are supported for both AWS and Azure environments. Depending on the cloud provider, Image Dedicated uses different technologies to establish private connectivity between Image Dedicated resources and user-managed cloud storage:

- AWS: Implemented through Amazon S3 Access Points to provide secure access to user-managed S3 buckets.
- Azure: Implemented through Azure Private Endpoints to establish a private connection to user-managed Azure Storage Accounts.

To read more about it and using them in Imagery Servers and Pro Machines refer to the Data Storage Accounts and Private Storage Connections Guide and Resources in the ArcGIS Image Dedicated site.

Notes:

- Azure private storage connections have no usage restrictions. They are supported across all versions of Imagery Servers, Pro Machines, and Processing and Analysis servers.

- AWS private storage connections have usage restrictions:
    - Supported only in Imagery Servers version ≥ 12.1
    - Supported in all Pro Machines
    - Not supported yet for Processing and Analysis services

Concepts

ArcGIS Image Dedicated Website Overview

You can register your ArcGIS Online organization in the ArcGIS Image Dedicated website to use ArcGIS Image Dedicated. ArcGIS Online identities are used to sign in and access the ArcGIS Image Dedicated website. To create or modify services, you will need to have an ArcGIS Online account with a Creator user type license.

Registration must be completed to log in to the website. Once registered, the organization administrators can perform management tasks, such as creating, updating, and deleting dynamic imagery, tiled imagery, RDBMS, Private Storage Connections and Data Storage Accounts and Pro Machines. Administrators can download the usage reports, which help them to better understand the resource consumption and costs. If users are assigned appropriate privileges to access defined items in their ArcGIS Online account, they can manage image services, perform processing and analysis tasks, create Pro Machines, and view the properties of the servers and RDBMS created in their organization. All users can download the Geoprocessing Tools and documentation.

What is an Image Service?

Image services provide a way to share raster and imagery data using configurable web services. With tiled imagery services, you can serve pre-processed imagery where the rendering and further processing are performed on the client. With dynamic image services, you can share imagery collections and perform server-side dynamic mosaicking and on-the-fly processing.

ArcGIS Image Dedicated has a set of custom geoprocessing tools that allow you to publish and manage image services from imagery and raster data stored in cloud storage.

Set Up

The creation and management of image services can be performed from a web-connected machine with ArcGIS Pro 3.3 or higher. These machines can either be local or a Pro Machine set up within ArcGIS Image Dedicated. In case you are using private storage connections in your workflows, these can be performed from a Pro Machine set up within ArcGIS Image Dedicated alone.

When using a local machine, MDTools must be downloaded and installed from the MDCS Github repo before you publish image services. MDTools is a command-line tool that simplifies common management tasks when working with rasters in a mosaic dataset.

You also need to download the ArcGIS Image Dedicated Service Management geoprocessing toolbox and add it to ArcGIS Pro to create and manage image services. The toolbox can be downloaded from Downloads tile on the ArcGIS Image Dedicated website.

To connect to an Enterprise Geodatabase created as part of the RDBMS and to author mosaic datasets, you need a Pro machine deployed in the same region as your RDBMS. Once connected to the Pro machine, you can access the Enterprise Geodatabase , author mosaic datasets, and publish them as dynamic imagery to your dynamic imagery server.

Dynamic Imagery

Overview

Dynamic imagery services can be used to publish rasters directly from cloud storage or from image definitions that contain mosaic datasets built from rasters. You can also publish tile caches, cloud raster formats, raster tiles, vector tiles, 3D tiles and scene layers as tile services if you have a dynamic imagery subscription.

To enable users to publish rasters—either directly to cloud storage or as an image service—ArcGIS Image Dedicated provides dedicated imagery servers that are managed by Esri and deployed in the same cloud and region as your preferred cloud storage. With a dynamic imagery subscription, you can publish both dynamic and tiled imagery services.

Tiled Imagery

Overview

Tiled imagery services can be used to publish tile caches, cloud raster formats, raster tiles, vector tiles, 3D tiles and scene layers as services directly from cloud storage. Datasets hosted in cloud storage can be accessed to create and update tile services.

To enable users to publish datasets stored in cloud storage, ArcGIS Image Dedicated provides dedicated tiled imagery servers to organizations with a tiled imagery subscription. These servers are managed by Esri and can be deployed in the same cloud and region as your cloud storage of interest. The tiled imagery server enables you to publish tile caches, raster tiles, CRF datasets as tile services, scene layers as scene services, 3D tiles as 3D tile services and vector tiles as vector tile services.

Note: Tiled imagery servers cannot serve dynamic imagery services.

Manage Cloud Stores in an Imagery Server

Imagery Servers can access (read) input data stored in cloud storage from:

- Data storage account specified during server creation.
- Cloud storage connection file (.acs file) provided during the service creation.

Because cloud storage credentials (e.g., access keys) may expire or be rotated, they must be updated periodically. The update method depends on how the cloud storage was configured with the server.

Data Storage Account

When an Imagery Server is created, you can specify a data storage account. This registers a default cloud store named aid_default in the server. The services on the server can access the associated cloud storage.

Updating Credentials of the default cloud store

AWS:

If the access keys of the data storage account are rotated:

- Sign in to the ArcGIS Image Dedicated site.
- Navigate to the Imagery Servers page.
- Locate the server whose credentials need to be refreshed.
- Click More Settings (three-dots icon) next to the server.
- Select Refresh Credentials.

The credentials will be updated automatically.

Azure:

If the security credentials of the data storage account are expired or rotated:

- Sign in to the ArcGIS Image Dedicated site.
- Create a new data storage account with the required details.
- Navigate to the Imagery Servers page.
- Locate the server whose credentials need to be updated.
- Click More Settings (three-dots icon) next to the server.
- Select Update Server Properties.
- Update the Data Storage Account Alias in the form and submit the job.

The credentials will be updated.

Cloud Storage Outside the Data Storage Account

If input data resides in a cloud storage location that is not part of the data storage account, credentials must be updated at the service level using an ACS File.
In this case, the update is performed using an Update Service action.

Update Credentials Using an ACS File

- Prepare an updated cloud storage connection file (ACS file) containing the new credentials.
- Create a new Image Definition (.zmd) with cloud storage connection file copied.
- Open the Manage Dynamic Services geoprocessing tool.
- Choose Update Service.
- Select the existing image service.
- Provide the updated Image Definition File.
- Submit the Job.

The service now uses new credentials contained in the ACS file.

Note: Updating credentials by submitting an Update Service request with a new ACS file must be repeated for each service using the cloud storage, which can be time-consuming in large environments.

Cloud Store Management Actions

To simplify cloud storage credential management, ArcGIS Image Dedicated provides support for registering and managing cloud stores directly on the Imagery Server using the user interface.

Benefits of registering the cloud store on the server:

- Centralized credential management.
- Credential updates apply to all services that reference the cloud store.
- There is no need to copy or update ACS files for individual services.
- Service-level updates are not required when credentials are rotated.

As part of cloud store management, you can perform the following actions:

Register a Cloud Store using a Cloud Storage Connection File (.acs).

Steps:

- Navigate to the Dynamic Imagery or Tiled Imagery page.
- Click the Server Alias.
- On the Server Details page, open the Cloud Stores tab.
- Click + Register a Cloud Store.
- Provide the Cloud Store Name and provide the Cloud Storage Connection File (.acs).
- Submit the job.
- Monitor the job status from the View Job Statuses table.

On successful completion, the cloud store is registered with the server.

Validate a Cloud Store

Validation confirms that the registered cloud store is accessible and configured correctly.

Steps:

- Open the Cloud Stores tab.
- Click More options (three-dots icon) next to the cloud store.
- Select Validate.

A validation job is submitted. You can monitor its status from the Job Status table.

Update Cloud Store

If validation fails or cloud storage credentials have expired or been rotated, you can update the cloud store credentials.

Steps:

- In the Cloud Stores tab, click More options next to the cloud store.
- Select Update.
- Provide a new Cloud Storage Connection File (.acs) containing valid credentials.
- Submit the job.

The cloud store credentials are updated, and all services referencing the cloud store automatically use the new credentials.

Note:

- The default cloud store (aid_default) cannot be updated using this workflow. To update it, refer to the section Data Storage Account.
- Uploaded ACS files are deleted after register/update operations and are not stored.

Usage Considerations and Limitations

- Supported providers: AWS and Microsoft Azure.
- Ensure the Region in the ACS file is valid.
- Cloud Storage Connection Files without any provider options are supported. If a provider option is specified, then AWS_NO_SIGN_REQUEST, AZURE_NO_SIGN_REQUEST, CPL_AZURE_ENDPOINT and AZURE_STORAGE_SAS_TOKEN are currently supported.
- After performing a Register or Update Cloud Store operation, the service might need to be stopped and started to reflect the changes. If the service is working as expected, restarting is not necessary.
- Changing the S3 bucket name in AWS, or the account name / container name in Azure cloud stores, is not supported. The Update Cloud Store operation only updates credentials. To use a different bucket, account, or container, a new cloud store must be created.
- Cloud Storage Connection Files created with public buckets or containers are supported in Imagery Server > 11.4.
    - Validation may fail with "Access Key ID Cannot be blank". This is a known limitation; the cloud store is still usable.
    - In Imagery Server 12.0, validation succeeds.

Service Management Toolbox

The Service Management toolbox contains tools to create and manage image services.

Tool	Description
Create Image Definition	Creates an image definition file used to manage dynamic or tile services.

Manage Dynamic Services	Creates and manages dynamic image services.
Manage Tile Services	Creates and manages tile image services.
Get Job Status	Allows to receive and review the status of jobs submitted by various tools.
Set System Service Settings	Helps in configuring the instance settings of the Raster Rendering Service and the Tile Services Shared Instance Pool.

Create Image Definition Tool

Overview

The Create Image Definition Geoprocessing tool aggregates and zips all the relevant data needed to create an image service into a. zmd extension file, otherwise known as the Image definition file. The file is saved locally on your computer and is used to publish dynamic and tiled image services.

This is the first tool used in the publishing workflow. This tool needs to be run whenever a compilation is ready for publishing or has been changed and an update needs to be published.

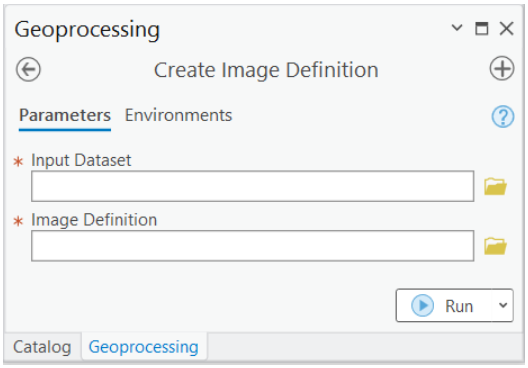

Parameters

Label	Explanation	Data Type
Input Dataset	The source path of a raster dataset, mosaic dataset, tile cache, raster tiles, vector tiles, or scene layer. The following are source path requirements: • **Raster Dataset**—The source path to the raster dataset should be a cloud path. • **Mosaic Dataset**—The source paths to the rasters added in the mosaic dataset should be cloud paths. If the mosaic dataset is made from raster proxies, the raster paths should be embedded. It can be from a File geodatabase or an Enterprise geodatabase. • **Tile Cache**—The source path to the tile cache should be a cloud path. The tile cache folder name must be in either of the formats "servicename" or "foldername_servicename." The "servicename" represents the tile service name and the "foldername" represents the server folder name.	File

	For example, if the tile cache name is "cities," it will be published as a service called cities in the root folder and the tile cache "world_cities" will be published as a service with the name "cities" in the folder with the name "world". If the tile cache is generated from an elevation dataset using the Elevation tiling scheme, which is a requirement for publishing it as a tile service and using it as an elevation layer in ArcGIS, it should be named in the format 'servicename_ImageServer'. This ensures that the tile cache folder name (excluding '_ImageServer') is published as the tile service name in the root folder. To specify a subfolder, the tile cache must be copied to a folder, the name of which will be utilized as the folder name. This should follow the format: '/arcgiscache/foldername/servicename_ImageServer'. • **Scene Layer, Raster Tiles, Vector Tiles, 3D Tiles**—The source path to these datasets should be a cloud path.	
Image Definition	The folder location and the name of the image definition being created. When storing the image definition in a folder, do not add a file extension to the name of the image definition because it will be added automatically.	File
Copy Cloud Storage Connection File	This parameter gets enabled when the input dataset used is a mosaic dataset. It determines whether the cloud storage connection file (.acs) should be copied into the image definition. If unchecked, the connection file will not be copied. The input dataset used to create the image definition should reside in a cloud storage (S3 Bucket or Azure Container) that was used during the creation of the dynamic imagery server. Unless using a mosaic dataset from a file geodatabase, the .acs file is always copied to the image definition.	Boolean

Get Job Status

Overview

The Get Job Status Geoprocessing tool allows you to review the status of jobs submitted by various tools. Each time you submit a job, a unique Job ID is returned. This is used as input to the Get Job Status tool.

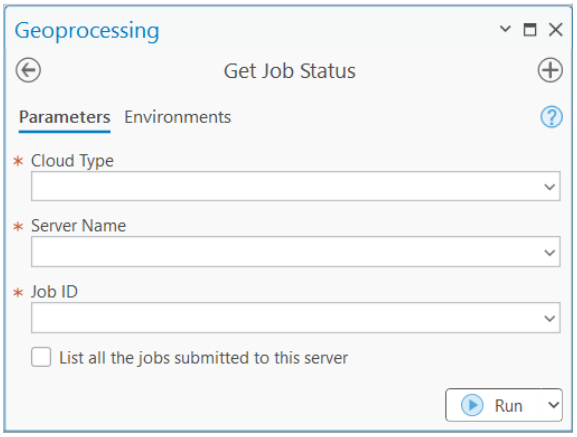

Parameters

Label	Explanation	Data Type
Cloud Type	The cloud type where the server is deployed.	String
Server Name	The server alias where the job was created.	String
Job ID	The Job ID of which the status is to be determined.	String
List all the jobs submitted to this server	Lists all the jobs submitted by other users within your organization on this server.	Boolean

Manage Dynamic Imagery

Overview

You can manage dynamic imagery using the Manage Dynamic Services Geoprocessing tool, which is a part of the ArcGIS Image Dedicated Service Management geoprocessing toolbox.

Using this tool, you can:

- Create, update, delete, start/stop, or refresh a dynamic image service.
- Set or revoke service permissions if you are the data creator.

The input for Manage Dynamic Services is image definition files that embed mosaic datasets or reference raster datasets.

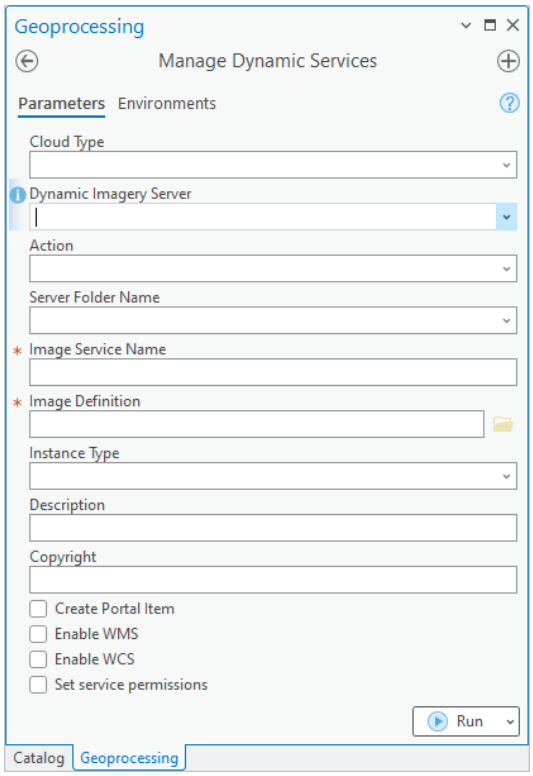

Parameters

Label	Explanation	Data Type
Cloud Type	The cloud type where the dynamic imagery server is deployed.	String

Label	Explanation	Data Type
Dynamic Imagery Server	The server alias where the image service is to be created or modified.	String
Action	The action that you would like to perform to manage the dynamic image service. The following actions are available: • **Create Service**—Creates an image service on the selected server. • **Update Service**—Updates an image service on the selected server. • **Delete Service**—Deletes an image service from the selected server. • **Start/Stop Service**—Starts or stops an image service on the selected server. • **Set Service Permission**—Sets or revokes service permissions of an image service on the selected server. • **Refresh Service**—Refreshes the image service on the selected server.	String
Server Folder Name (Optional)	The name of the folder to be used by the dynamic image service. The following actions are available: • **Create Service**—Select an existing folder or enter a new folder name to create a folder. • **Update Service/Delete Service/Refresh Service/Start Service/Stop Service/Set Service Permission**—Select the folder name in which the service is present.	String
Image Service Name	The name for the image service.	String
Image Definition	The image definition file (.zmd) created by the "Create Image Definition" tool. e.g., C:\Projects\imagedefinitions\sample.zmd Image definitions of size greater than 2 GB can be uploaded to an Amazon S3 bucket or Azure container. The cloud paths to these files are to be entered manually. The cloud storage used to store the image definition file can be public or secure. If the storage is secured, ensure that server has read access to it. The following are the different cloud storage paths that are supported: • Cloud storage connection • S3 URI • VSI file handler (/vsis3)	File
Instance Type (Optional)	The instance type that will be used for the image service. Requests made to a service are handled by processes running on the server. Each process uses server memory to run, so it is important to limit the total number of processes. The following actions are available:	String

Label	Explanation	Data Type
	• **Shared instances**—Conserves memory usage by using a pool of processes shared by multiple services. The disadvantage of this approach is that the first request to a service may take a few more seconds to start up. Using a shared instance is recommended for larger numbers of services that each receive infrequent requests. • **Dedicated instances**—Uses a reserved pool of processes specific to the service. This uses more memory but is ideal for services that receive constant or compute-intensive requests.	
Minimum Instances	Specifies the minimum number of dedicated instances that can run for the service at any given time, even when the service is not being used.	String
Maximum Instances	Specifies the maximum number of dedicated instances that can run for the service at any given time, even when the service is not being used.	String
Description (Optional)	A description of the image service to be created or updated.	String
Copyright (Optional)	Copyright text for the image service to be created or updated.	String
Users to give or revoke permission (Optional)	Lists the users and their permission details. The following actions are available: • **User**—Specifies the username for which permissions will be applied. • **Permission detail**—Specifies whether permission is enabled or disabled.	String
Change service status (Optional)	Shows the status of the selected service. Change the value to update the status of the service. The following actions are available: • **STARTED**—The service has started. • **STOPPED**—The service is stopped.	String
Create Portal Item (Optional)	Creates a portal item for the created image service URL. If selected, a portal item will be created, and the service will be referenced using this item. If left unchecked, only the service endpoint will be created. You can register the service endpoint to ArcGIS Online or ArcGIS Enterprise portals.	Boolean
Enable WMS (Optional)	Enables WMS capabilities and update WMS properties.	Boolean
Enable WCS (Optional)	Enables WCS capabilities and update WCS properties.	Boolean
Set service permissions (Optional)	Set permissions required to access the image service. If permissions are not set, the service will be accessible by any user with the service URL.	Boolean

Note: If the Manage Dynamic Services tool is not listing any servers or updating the server list despite having server access, delete the JSON file created with your username in the di_profiles and ti_profiles folders in the toolbox location. Re-run the tool, and the issue should be resolved.

Manage Tiled Imagery

Overview

You can manage tiled imagery using the Manage Tile Services Geoprocessing tool, which is a part of the ArcGIS Image Dedicated Service Management geoprocessing toolbox. With this tool, you can create, update, delete, or start/stop a tiled image service. You can also set or revoke service permissions if you are the data creator.

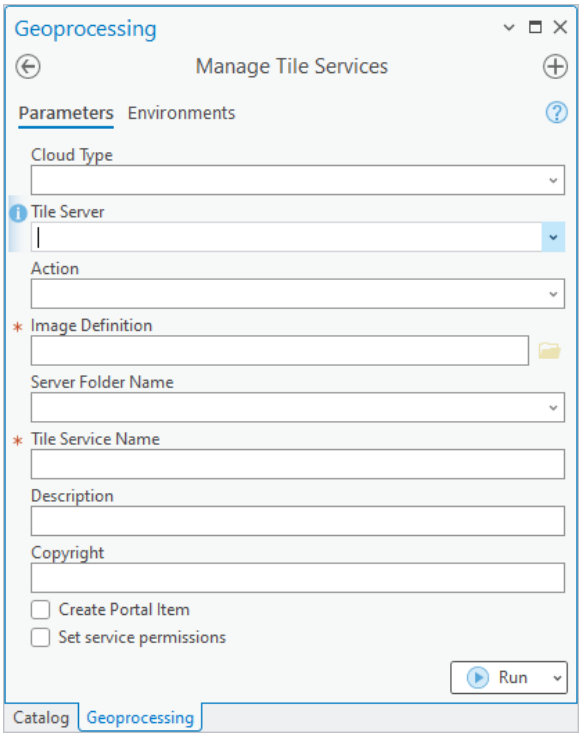

Parameters

Label	Explanation	Data Type
Cloud Type	The cloud type where the tile server is deployed.	String
Tile Server	The server alias where the tile service is to be created or modified.	String
Action	The action you would like to perform to manage the tile service. The following actions are available: • **Create Service**—Creates a tile service on the selected server. • **Update Service**—Updates a tile service on the selected server. • **Delete Service**—Deletes a tile service from the selected server. • **Start/Stop Service**—Starts or stops a tile service on the selected server. • **Set Service Permission**—Sets or revokes service permissions of a tile service on the selected server. • **Clear REST Cache**–Clears the REST cache.	String

Label	Explanation	Data Type
Image Definition	The image definition file (.zmd) created by the "Create Image Definition" tool. (e.g., C:\Projects\imagedefinitions\sample.zmd)	File
Server Folder Name (Optional)	The server folder name where the tile service will be placed. The following actions are available: • **Create Service**—Select an existing folder or enter a new folder name to create a folder. • **Update Service/Delete Service/Start Service/Stop Service/Set Service Permission/Clear REST Cache**—Select the folder name where the service is located. When creating a service with a tile cache, the server folder name gets populated automatically using the provided image definition.	String
Tile Service Name	The name of the tile service or scene layer. If using tile cache, the tile service name populates automatically using the provided image definition.	String
Instance Type (Optional)	The instance type that will be used for the tiled image service (CRF as input dataset). Requests made to a service are handled by processes running on the server. Each process uses server memory to run, so it is important to limit the total number of processes. The following actions are available: • **Shared instances**—Conserves memory usage by using a pool of processes shared by multiple services. The disadvantage of this approach is that the first request to a service may take a few more seconds to start up. Using a shared instance is recommended for larger numbers of services that each receive infrequent requests. • **Dedicated instances**—Uses a reserved pool of processes specific to the service. This uses more memory but is ideal for services that receive constant or compute-intensive requests.	String
Minimum Instances	Specifies the minimum number of dedicated instances that can run for the service at any given time, even when the service is not being used.	String
Maximum Instances	Specifies the maximum number of dedicated instances that can run for the service at any given time, even when the service is not being used.	String
Description (Optional)	If using tile cache, scene layer, raster tiles, or vector tiles description must be given and updated in the root.json manually.	String
Copyright (Optional)	If using tile cache, scene layer, raster tiles, or vector tiles, the copyright text must be given and updated in the root.json manually.	String
Users to give or revoke permission (Optional)	Lists the users and their permission details. The following actions are available: • **User**—Specifies the username for which permissions will be applied. • **Permission detail**—Specifies whether permission is enabled or disabled.	String
Change service status (Optional)	Shows the status of the selected service. Change the value to update the status of the service. The following actions are available:	String

Label	Explanation	Data Type
	• **STARTED**—The service is started. • **STOPPED**—The service is stopped.	
Create Portal Item (Optional)	If checked, creates a portal item for the created tile service URL.	Boolean
Set service permissions (Optional)	If checked, set permissions required to access the tile service. If unselected, the service will be accessible by any user with the service URL.	Boolean

Set System Service Settings

Overview

You can configure the instance settings of the Raster Rendering Service and Tile Services Shared Instance Pool using the Set System Service Settings Geoprocessing tool, which is a part of the ArcGIS Image Dedicated Service Management geoprocessing toolbox.

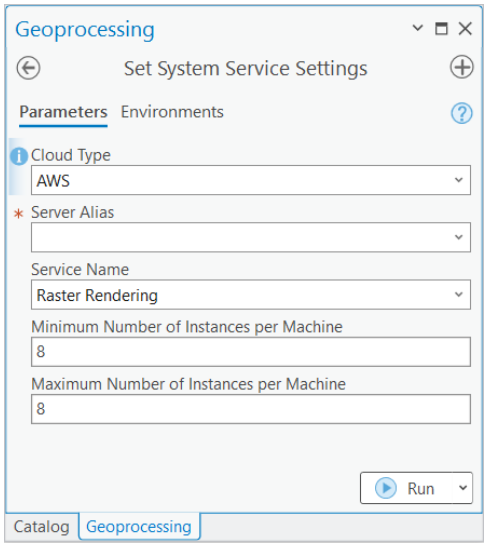

Parameters

Label	Explanation	Data Type
Cloud Type	The cloud type where the server is deployed.	String
Server Alias	The server alias whose instance settings are to be updated.	String
Service Name	The service name whose instance settings are to be updated. The following services are available: • **Raster Rendering**— This service is a multi-tenant raster rendering provider. Shared instance image services make use of this service, which allows many services to share a pool of compute processes. This enables a large number of different services to be served from more limited resources.	String

	• **Tile Services Shared Instance Pool**— Tile services (except Tiled Image Services) published to the tiled imagery server use this shared instance pool. Shared instances conserve memory usage by pooling active instances for use by multiple services. It can reduce the overall number of instances you need to run to meet your service traffic needs. This makes it easier to strike the optimal balance between performance and memory usage.	
Minimum Number of Instances per Machine	The minimum number of dedicated instances per machine.	String
Maximum Number of Instances per Machine	The minimum number of dedicated instances per machine.	String
Number of Shared Instances in the Pool	The number of shared instances in a pool per machine.	String

Processing and Analysis

Overview

If you have subscribed to Processing and Analysis within ArcGIS Image Dedicated, you have a set number of Image Dedicated credits that you can use to run analyses or image processing tools unique to Image Dedicated.

The Processing and Analysis services can be used to pre-process data to optimize it for publishing or to perform analysis on the imagery to extract information products. When jobs are submitted, Esri launches the appropriate cloud computing instances to process or analyze the imagery in the same cloud region where the source imagery resides. This ensures faster processing and significantly reduces egress costs.

Enabling or Disabling a Processing and Analysis Service Task

A Processing and Analysis service can encompass multiple tasks. To access the Processing and Analysis Tasks page, click on the 'Processing and Analysis' tab in the ArcGIS Image Dedicated Site.

Upon navigating to the Processing and Analysis Tasks page, you will find three tabs, Pro Machine, RDBMS and Geoprocessing Tasks. Click on any tab; all the tasks/activities relevant to each of them are listed. Further details about individual tasks can be accessed by clicking on the respective service/task.

In the Geoprocessing Tasks tab, once you identify a service/task that suits your requirements, click on it to explore the service. If it proves suitable, enable the task to commence submitting the jobs.

To enable a task, use the toggle next to it. Enabling a task prompts you to either create a new access item in your ArcGIS Online account or use an existing item to regulate access to the processing and analysis task, as well as job submissions. After enabling a task, you can view details such as its current access item, your access status, and options to view or update the current access item. You can view the list of Job IDs and Status; you can subscribe or unsubscribe from email notifications.

Each Processing and Analysis job submission consumes ArcGIS Image Dedicated credits. After completing your tasks, consider disabling a specific task to prevent accidental job submissions and credit consumption.

Output to Cloud Storage

The output from Processing and Analysis can be output to a defined cloud storage that should be in the same region as the processing to reduce egress costs.

Publish Results

You can also publish your results to ArcGIS Online and ArcGIS Enterprise and create associated portal items. You can set this up when publishing dynamic and tiled image services using the Manage Dynamic Imagery or Manage Tile Imagery geoprocessing tools. The image service will be referenced using this portal item in your ArcGIS organization.

Items are classified as either web content or desktop content. Desktop content is primarily consumed by desktop applications and must be published as geospatial web services before they can be widely consumed over the web. Web content can be actively consumed across all client platforms, including web browsers, mobile devices, tablets, and desktops.

You can use a portal to store and share content among users who can directly consume the services or download the content if applicable. If you do not create a portal item, only the service endpoint will be created. However, you can also register the service endpoint to ArcGIS Online or ArcGIS Enterprise. To do this, you can use the web interface of ArcGIS Online or ArcGIS Enterprise to create items that reference services by providing the REST endpoint and credentials from ArcGIS Image Dedicated servers. If you don't define the credentials, the item would become a link to the service. However, if you do provide credentials, then the portal becomes a proxy server to the item. In this case, all requests will go to the portal, which will then append the required credentials and send the request. You will also receive the response through the portal. This additional step ensures suitable security; however, it does negatively impact performance, especially if the server is in a different region than the portal.

Managed RDBMS

Overview

As part of the Processing and Analysis subscription, you can also leverage managed RDBMS. It provides a fully managed relational database deployed alongside the Dynamic Imagery Server in the cloud. With Managed RDBMS, users do not need to provision or maintain any database infrastructure themselves, enabling a streamlined, cloud-native experience for enterprise imagery management. An Enterprise Geodatabase is pre-created and made available as a ready-to-use enterprise geodatabase connection file (.sde), allowing immediate use in building mosaic datasets and publishing them as dynamic image services through the Dynamic Imagery Server. With the database server hosted in the same cloud environment as the Imagery Server and Pro Machines, you will benefit from seamless integration, efficient data access, and reduced complexity in dynamic imagery workflows. The RDBMS and its access are managed through the ArcGIS Image Dedicated website. Access to the database is provided using an enterprise geodatabase file. To use it, you need to connect to a Pro Machine via Remote Desktop.

Create an RDBMS

To create a new RDBMS, you must have a dynamic imagery server (version 11.4 or higher) up and running. Once the imagery server is in place, ensure that you have sufficient Image Dedicated Credits purchased through a Processing and Analysis subscription. Credits will apply based on the selected term length of the RDBMS. Execute these steps to create an RDBMS.

Note: Image Dedicated Credits consumption is based on the term length selected for creating the RDBMS.

Term Length	Image Dedicated Credits
3 Months	4320
6 Months	8640
12 Months	17520

Currently, only PostgreSQL is available as the RDBMS type.

Once you submit your request form, the RDBMS will be created using the name, database server alias, and specifications you provided in the request form. This database will then be enabled as an enterprise geodatabase, completing the RDBMS creation process.

It typically takes up to 15 minutes for the RDBMS to be created and listed as an active RDBMS on the Image Dedicated website. It is recommended to use a unique name each time you create a new RDBMS to differentiate between database servers.

Access an RDBMS

The database server created as part of the RDBMS is a private resource and is not accessible over the internet. It can only be accessed from a Pro machine and the configured dynamic imagery server. To access an RDBMS created in a specific region, both the Pro machine and the dynamic imagery server must also be located in the same region. Cross-region access is not supported.

Once the database server status changes to “started” in the list of active RDBMS, you can download the enterprise geodatabase connection file by clicking the three dots next to your database server. This connection file can then be copied to a Pro machine with access to the RDBMS. You can add the connection file in ArcGIS Pro to connect to your database.

Please note that you cannot change the username or password of your database, nor can you create additional users. You will receive a single-user enterprise geodatabase.

Manage RDBMS Access

By default, the dynamic imagery server used to create the RDBMS will have read access to the database. Access for other resources, such as additional running imagery servers or Pro machines, must be explicitly granted as needed. Access management varies depending on the cloud provider.

In Azure, all imagery servers and Pro machines running in your account and within the same region as the RDBMS will automatically have access.

In AWS, you must explicitly enable access for additional resources, such as dynamic imagery servers or pro machines, even if they are in the same region as the RDBMS. Similarly, access can be disabled at any time.

To enable or disable access to the RDBMS:

- For a Pro Machine – Click the three dots next to the Machine alias, open the Manage RDBMS Access form, and select the RDBMS to enable or disable access.
- For a Dynamic Imagery Server - Click the three dots next to the Server alias, open the Manage RDBMS Access form, and select the RDBMS to enable or disable access.

Note: Only RDBMS located in the same region as the Imagery Server or Pro Machine will be available for access requests.

Note: In the AWS environment, if you are unable to connect to the database using the enterprise geodatabase connection file, check if your Pro machine has access enabled. If not, you must enable access before connecting to the database.

Note: Be cautious when granting access to the RDBMS for Pro machines, as data can be deleted once connected. Only administrators can access the Manage RDBMS Access form and enable or disable RDBMS access for imagery servers and Pro machines.

Use an RDBMS

Once everything is set up, you can copy the enterprise geodatabase connection file to your Pro machine. Follow these steps to access the database:

- Create a project in ArcGIS Pro.
- In the Catalog pane, right-click on Databases and select Add Database.
- Browse to the location where you copied the enterprise geodatabase connection file and select it.
- The enterprise geodatabase file will be added to the catalog pane under Databases.
- Expand the enterprise geodatabase to author mosaic datasets.

To publish a mosaic dataset authored in an enterprise geodatabase to a dynamic imagery server, the imagery server must have read access to the RDBMS where the mosaic dataset resides. In the AWS environment, ensure that access to the RDBMS is enabled for the dynamic imagery server, as explained in the previous section. Otherwise, the service creation operation will fail because the server cannot access the database.

Extend the Term Length of a Running RDBMS

You can extend the term length of a running RDBMS by selecting an additional usage period.

1. Open the ArcGIS Image Dedicated website
    - Navigate to the Processing and Analysis page.

2. Open the RDBMS options

- Click the RDBMS tile.
- Locate the RDBMS you want to extend; click More options (…).
- Click Extend Term Length.

3. Select a term extension

Choose one of the available terms:

- 3 Months
- 6 Months
- 12 Months

4. Submit your request

- Confirm and submit the term extension request.
- The selected duration is added to the current term.
- Image Dedicated credits for the selected term length are deducted.

Your RDBMS will now continue running for the extended term without interruption.

Upgrade an RDBMS

The engine version of the database server created as part of the RDBMS can be upgraded to a higher version to ensure compatibility with newer imagery server requirements. This is typically required when upgrading an imagery server to a newer version that includes updated database prerequisites.

To upgrade the database server engine version, click the three dots next to the database server and open the Upgrade Database Server Version form. Specify the Reference Server (the only required input), which determines the target version for the database server upgrade; the remaining fields—Type of RDBMS, Current Database Version, and Database Version to Upgrade—are auto-populated based on the selected database server and Reference server.

Click Upgrade Database Server to initiate the process. This action submits a job to upgrade the database server engine version.

Note:

Upgrading the database server is an irreversible action and cannot be undone. The process will initiate a version upgrade to the selected version, and downtime is expected during this period. The duration of downtime may vary depending on the database size and the volume of stored data.

Important:

If you upgrade the RDBMS to meet the requirements of a newer imagery server version (for example, upgrading to support an Imagery Server 12.0 deployment), the upgraded database server may still

function with older imagery server versions (such as 11.5). However, the upgraded database version may not meet the officially supported and certified requirements for the older imagery server version.

It is recommended to verify database compatibility against the official system requirements documentation before performing such upgrades. For more details, refer to the ArcGIS Server requirements for using databases.

Pro Machines

Overview

As part of the Processing and Analysis subscription, you can also leverage virtual Pro Machines. A Pro Machine can be created in the same cloud and region as your imagery data to publish and manage image services and perform analysis. Pro Machines are managed on the ArcGIS Image Dedicated website and can be accessed using Remote Desktop.

You will need to bring your own ArcGIS Pro license (Basic, Standard, or Advanced) to use ArcGIS Pro on a Pro Machine. Licenses can also contain any extension, such as Image Analyst, which is often required in image analysis workflows.
ArcGIS Pro licenses are now primarily managed through ArcGIS Online accounts. Organizations can assign Named User licenses directly from their ArcGIS Online portal, allowing users to access ArcGIS Pro based on their assigned roles and permissions.

Create a Pro Machine

To create a new virtual Pro Machine, you must complete a machine request form, which is available on the ArcGIS Image Dedicated website.

Go to the **ArcGIS Image Dedicated** website.

Select the **Pro Machine** tile on the homepage.

Click **+ New Machine**. Fill out the required details as prompted: For example,

- **Cloud Provider**: Choose where your imagery is stored (e.g., AWS, Azure).
- **Designated Cloud Region:** Select the appropriate region.
- **Pro Version** – Select the ArcGIS Pro version used for deployment.
- **Instance Type**: Select the configuration that suits your performance needs.

The instance types that are available in the Pro Machine creation form correspond to AWS and Microsoft Azure cloud storage models. Six sizes, ranging from Small Central Processing Unit (CPU) to Large Graphics Processing Unit (GPU), each with different costs and increasing performance, are available.

For AWS, there are six instance types that correspond with six AWS model types associated with EC2 M5 instances, I3 instances, and G6 instances. For Pro 3.5, CPU instance types have been upgraded to use M6id, while earlier Pro versions will continue to use M5. Fallback GPU instances (g4dn.xlarge family) are listed alongside the primary G6 instances. These fallbacks are automatic options in regions where the

specified G6 instances are not available. Three of the instances are CPU-based and the other three are GPU-based. More information—such as the memory, storage, and processing type—is provided in the following table. Instance type and cost correspond with the Pro Machine request form. All other fields correspond with AWS model types and the associated information.

Note: The machine types and specifications may be different depending on availability.

Instance Type	Image Dedicated Credits (per hour)	AWS Model /Fallback GPU Instance	Processors	Memory	Instance Storage
Small_CPU	1	m5d.xlarge	4 (vCPU)	16 GiB	150 GiB
Small_CPU (Pro version 3.5)	1	m6id.xlarge	4 (vCPU)	16 GiB	237 GiB
Medium_CPU	1.5	m5d.2xlarge	8 (vCPU)	32 GiB	300 GiB
Medium_CPU (Pro version 3.5)	1.5	m6id.2xlarge	8 (vCPU)	32 GiB	474 GiB
Large_CPU	2	i3.2xlarge	8 (vCPU)	61 GiB	1.9 TB
Large_CPU (Pro version 3.5)	2	i7i.2xlarge	8 (vCPU)	64 GiB	1.7 TB
Small_GPU	2	g6.xlarge / g4dn.xlarge	4 (vCPU) + GPU L4 / 4 (vCPU) + GPU T4	16 GiB	250 GiB /125 GiB
Medium_GPU	3	g6.2xlarge / g4dn.2xlarge	8 (vCPU) + GPU L4 / 8 (vCPU) + GPU T4	32 GiB	450 GiB /225 GiB
Large_GPU	5	g6.4xlarge / g4dn.4xlarge	16 (vCPU) + GPU L4 / 16 (vCPU) + GPU T4	64 GiB	600 GiB /225 GiB

For Microsoft Azure, there are six instance types that correspond with Azure D-series, F-series, NC-series, and NV-series virtual machines. Three of the instances are CPU-based, and the other three are GPU-based. Fallback GPU instance (Standard_NC16as_T4_v3) is listed alongside the Large GPU instance. This fallback is automatic option in regions where the Standard_NV12ads_A10_v5 instances is not available. More information—such as the memory, storage, and processing type—is provided in the following table. Instance type and cost correspond with the Pro Machine request form. All other fields correspond with Azure model types and the associated information.

Instance Type	Image Dedicated Credits (per hour)	Azure Model/ Fallback GPU Instance	Logical Processors	Memory	Temp Disk Storage
Small_CPU	1	Standard_D4ds_v5	4 (vCPU)	16 GiB	150 GiB
Medium_CPU	1.5	Standard_F8s_v2	8 (vCPU)	16 GiB	64 GiB
Large_CPU	2	Standard_D8ds_v5	8 (vCPU)	32 GiB	330 GiB
Small_GPU	2	Standard_NC4as_T4_v3	4 (vCPU) + GPU T4	28 GiB	180 GiB
Medium_GPU	3	Standard_NC8as_T4_v3	8 (vCPU) + GPU T4	56 GiB	360 GiB
Large_GPU	5	Standard_NV12ads_A10_v5 / Standard_NC16as_T4_v3	12 (vCPU) + GPU 1/3X A10 /16 (vCPU) + + GPU T4	110 GiB /110 GiB	320 GiB

- **Apply Templates**– Select a Template from the available list. Templates contain pre-installed components for the Pro machine. The templates available are dependent on regions or instance types.

You can preconfigure the Pro Machine with a set of geoprocessing tools or datasets to customise them for your workflow. For the same, when creating/adding a new Pro Machine, select a template to use the pre-configuration (template data) for the Pro Machine. The respective ACS files, Imagery Workflows file, and ReadMe file on the template will be copied to the Pro Machine during start-up. The available template list differs based on the cloud type and cloud region.

Note: When you choose a Deep Learning template, all the necessary libraries for deep learning are automatically installed within the pro machine's environment.

- **Additional Disk Storage**: You can add additional disk storage to either a new Pro machine or to an existing Pro machine that is in Started Status. Select the Additional Disk Storage from the available list. The selected additional storage will be allocated in the Pro Machine. The cost for the additional disk storage will be Images Dedicated Credits as listed in the table below. Users will be charged monthly for the additional storage regardless of the machine status, with the amount depending on the selected storage size. To ensure users are aware of these ongoing costs, an alert message is displayed next to the status of an inactive machine.

Note: In the AWS environment, if you need to update additional disk storage on a Pro Machine, take note that there is a wait time of at least six hours between successive disk storage modifications. However, in the Azure environment, there is no specific waiting limit.

Additional Disk Storage	Image Dedicated Credits (per month)
250 GiB	52.5
500 GiB	105
750 GiB	157.5
1 TiB	210
1.5 TiB	315
2 TiB	420

- **Inbound access IP address preferences** – Enter an IP address. The RDP connection is enabled to the provided IP Address/es.

ArcGIS Image Dedicated allows a maximum of three IP addresses for a Pro Machine. The initial IP address will be automatically populated based on your connected network. If you intend to connect to the Pro Machine with different networks, additional IP addresses must be provided. To add IP address, click **+ Add Additional IP Address**. Additional IP addresses cannot be added when the machine is in the Stopped state. You may also add an optional note for each IP address to help identify and manage them more easily. The note fields are aligned with the IP address field, allowing you to add or edit notes for

each IP address as needed. You will not be able to connect to the Pro Machine if you try to access the machine with an unspecified IP address. The IP address that you selected when creating the Pro Machine will enable the Remote Desktop Protocol (RDP) connection and help you access the Pro Machine.

- **Pro Machine Alias** – Enter a name for the ArcGIS Pro Machine.
- **Additional Comments** – Add any specific information regarding the machine.

Once the request form is submitted, the Pro Machine will be created with the machine's alias and specifications provided. It will typically take up to 10 minutes for the Pro Machine to be listed as active on the Image Dedicated website. It is recommended that you use a unique name for every Pro Machine created.

Note: It is recommended that the Pro Machine be spun up in the same cloud region as the data storage to avoid latency and egress costs.

You can click on the **Machine alias** of the Pro Machine to list the machine details. You can modify and update to apply the changes. You are notified of the changes.

Access a Pro Machine

In earlier versions of ArcGIS Image Dedicated, users had to connect directly through the RDP port 3389, which often required administrators to whitelist the IP address and enable the RDP port 3389. For Pro Machines 3.4 and above, the user connects using Remote Desktop Gateway (default). However, the choice to connect using Remote Desktop file is still available to the user.

RD Gateway uses the HTTPS protocol to encrypt and secure the connection, encapsulating RDP within HTTPS and using SSL/TLS for encryption. This extra layer of security routes all traffic through port 443 (HTTPS), providing encryption and secure transport. With DNS over HTTPS, if the user must request whitelisting, they need to request domain whitelisting and not IP address whitelisting, making it a one-time task. Currently, the DNS provided is *.img.ArcGIS.com.

Once your Pro Machine is created, to connect through Remote Desktop use the credentials provided on the ArcGIS Image Dedicated website. To access these credentials, go to the ArcGIS Image Dedicated homepage, open the **Processing and Analysis** tab, and select the **Pro Machine** page. This page lists all active Pro Machines along with their associated credentials. Machines are grouped by cloud provider, and the default view shows machines for the provider you initially selected. To view machines hosted on other cloud providers, use the **Cloud Provider** drop-down menu and select the one you want.

Select to copy the IP address/ Gateway Server, username, and password. Alternatively, you can download the RDP file/ RDP file (Remote Desktop Gateway) containing the credentials using **Retrieve Credentials** and use the same to connect to the Pro Machine. Alternatively, you can also connect using the Microsoft Terminal Services Client (MSTSC) command. This is because the Pro Machine is virtual and uses a remote desktop in the cloud region specified when completing a machine request form.

If you are unable to connect to the Pro Machine using the RDP file or MSTSC command, ensure the machine's IP address matches the IP address added as the Inbound access IP address for the machine. Otherwise, update the IP address for the machine to connect to the virtual Pro Machine.

Using a Pro Machine

Data Management

Data on the D Drive—which is an ephemeral drive if created in Azure—and the Z drive—which is an ephemeral drive if created in AWS or Azure—will be lost if the Pro Machine is stopped or impaired. The D and Z are basically the same drives. Ideally, refrain from using the ephemeral Z drive or the D drive to store data. These drives are fast drives and are recommended to be used for temporarily storing data to speed up processing. If you have not opted for additional disk storage while creating Pro Machine, it is optimal to copy pre-processed data to the D and Z drives. Perform processing with the output directed to the same drive. Alternatively, copy it to the additional disk storage (if any). In such scenarios, the result would need to be copied back to the E drive, additional disk storage, or copied back to the cloud (or other storage).

Ensure any important data is backed up frequently in the E drive or additional disk storage. Any new projects created in a Pro Machine will be in the default directory E:\Projects folder. Do not delete or rename the folders in the E drive. The folder structure maintained in the E drive contains a collection of folders to organize different file types. The folder structure is as follows:

- **ACSFiles**—This folder contains the pre-created ACS connection files of the cloud and region-specific public datasets available at Registry of Open Data on AWS and Microsoft's Planetary Computer Data Catalog.
- **Data**—This folder is used to store your data such as raster datasets. Other project-related resources can also be stored in this folder.
- **ImageryWorkflows**—This folder contains the workflows to build mosaic datasets and manage, analyze, or use imagery and rasters. The workflows available in this folder change depending on the template selected when creating a Pro Machine.
- **Projects**—This is the default location for the projects that are created in ArcGIS Pro.
- **PythonNotebooks**—This folder contains the ArcGIS Notebooks.

Data Access

Although users can access their own cloud storage in a virtual Pro Machine, discussed in Cloud Storage, the Pro Machine also comes with open imagery data from public AWS buckets and Azure containers. ArcGIS Image Dedicated supports the open datasets that cloud providers offer as pre-made ACS connection files. This saves the user time so they can use open datasets for image processing and analysis tasks. The ACS files automatically populate the present data from these open-source buckets. The ACS files are found on the E drive of the Pro Machine.

Stop a Pro Machine

When you finish working on a Pro Machine, restart or shut down the machine using the machine's power button. It is a good practice to Stop the active machine when not in use to prevent additional costs. However, note that the data in the ephemeral (D and Z) drives will be lost each time the machine

is shut down. You can also Start and Stop the Pro Machine from the Image Dedicated website ->Processing and Analysis ->Pro Machine->Active Pro Machines list. After you submit the stop request, the machine stops immediately and is powered off. The machine status updates to Stopped to indicate that it is no longer running or incurring charges. Machine credentials are regenerated each time the machine is started or stopped. As a result, you must download a new RDP file to connect using the updated credentials. This requirement does not apply when connecting through the Remote Desktop Gateway, where re-downloading credentials is not necessary.

Pro Timer

When you connect to a Pro Machine, the Pro Timer will automatically start with a default timer set for eight hours and run during the entire session. Once the 8-hour timeframe is complete, the machine will be shut down automatically. You can reset or extend the timer depending on your requirements. The timer can be extended to a maximum of 720 hours (or 30 days).

Other than assisting in setting timers to automatically shut down your Pro Machine, you can also back up your data to cloud storage (AWS S3 bucket or Azure container) with the Pro Timer. This feature ensures your data will be saved to cloud storage before the machine shuts down.

The Pro Timer also displays a caution to the users prompting about the Windows updates enabled on the Pro Machines. Even though users are notified before the actual updates, we recommend proceeding with caution, as some updates may require restarting the machine, potentially interrupting operations and processes. To avoid the interruptions, users can choose to temporarily pause Windows updates by navigating to Windows Update from System Settings.

Reset Password

The Pro machine password expires after 42 days. You will be alerted with a message when attempting to log in using the expired password. To reset your Pro machine's password:

- Navigate to the Pro Machine Alias for which you wish to reset the password.
- Scroll down to the Password section. Click "Reset". Your password will be reset within 5 minutes and be notified.
- Copy the new password from the "Retrieve Credentials" page. This password can be used to log in to your Pro machine.

Note: The password for a Pro machine in the Updating or Stopped state cannot be reset. Ensure the Pro Machine is in the Started state to reset the password.

Delete a Pro Machine

You can delete existing Pro Machines by selecting the machine from the Active Pro Machines list. When a machine is deleted, the data stored on all the drives is deleted. Click **Machine Alias**. In the response window and scroll down to the **Delete Protection** section. Ensure the option **Prevent this machine from being accidentally deleted** is unchecked. Click on **Delete Machine**. The users are notified of the changes. **Note:** If the option **Prevent this machine from being accidentally deleted** is selected, you cannot delete the machine. Similarly, if you would like to protect your Pro Machine from accidentally getting deleted as an extra safety precaution, it is a good practice to enable the option.

Before deleting a Pro Machine, ensure to de-authorize your single or concurrent use licenses of ArcGIS Pro, along with the extensions. Otherwise, the licenses will still be in an authorized state, and you may not be able to authorize them on other machines without resetting them.

Administration

Email Notifications

Email notifications are sent to inform you about the management activities in your ArcGIS Image Dedicated account, such as the status of new dynamic imagery server or tiled imagery server creation jobs submitted, the status of processing and analysis jobs submitted (whether they have succeeded or failed), and if ArcGIS Image Dedicated Credits have gone negative or reached an exhaustion level. These emails will be sent to the administrator who is marked as the point of contact when subscribing to ArcGIS Image Dedicated.

There is a set of email notifications to which you can subscribe or unsubscribe. For example, Processing and Analysis.

However, there is another set of email notifications that cannot be unsubscribed from, such as notifications to inform you about the server status whenever you submit a new server request or notifications about ArcGIS Image Dedicated Credits going low or negative.

Usage Reports

The usage reports of the Dynamic Imagery Server, Tile Server, Pro Machine, and other Processing and Analysis services can be generated from the ArcGIS Image Dedicated website. These reports will help you understand the usage of these resources and, in turn, provide the costs associated with the resources used. To access usage reports, select the Admin->Reports, select the service you would like to get a report for, and complete the form with your specifications as prompted.

You can also view your Account Summary (how many Image Dedicated credits are left) using the ArcGIS Image Dedicated website. To view your credit usage, click on Admin->Account Summary.

Control Access Items

ArcGIS Online allows users to store and share a variety of items. Each item has a type, such as web maps, feature services, or web applications, as well as a set of type keywords that provide additional information about the characteristics of that item type. ArcGIS Online can be used to store and share this content among users of an organization or multiple organizations through groups. Each item has a unique identifier and a known URL independent of the user owning the item.

ArcGIS Image Dedicated uses access items to control who can access management and processing capabilities. Access items are created when an ArcGIS Online user with an administrator or publisher role subscribes to Image Dedicated and creates a dynamic imagery server or tiled imagery server. Whereas, in the case of processing and analysis services (such as creating a Pro Machine or performing Processing and Analysis jobs), the administrator can create a new access item or use an existing access item in ArcGIS

Online when enabling a Processing and Analysis Task from the Processing and Analysis chiclet in the ArcGIS Image Dedicated Site.

As an administrator, you can decide/provide access to the appropriate items to users and publishers in your organization for all management and processing capabilities.

There are three types of access items in Image Dedicated for authentication. They are:

- Admin Access Item
- Server Access Item
- P&A Access Item

Admin Access Item

The Admin Access Item allows a user to administer ArcGIS Image Dedicated for their ArcGIS Online organization. Refer to Registration section of the document for more details.

Server Access Item

The Server Access Item controls access to servers that manage dynamic and tiled imagery. A user with access to this item can manage image services, update server properties, and delete servers. This access item is created when a dynamic or tiled imagery server is created.

To create a Server Access Item, you must select Create a Server Access Item on your ArcGIS Online account when filling out the server request form. The administrator can share the access item with other members of the organization who would like to manage and publish image services to the dynamic or tiled imagery server.

Complete the following steps to create a server access item:

1. Log in to ArcGIS Image Dedicated site.
2. Click the Dynamic Imagery or the Tiled Imagery tile, depending on which server you want to create (and your subscription).
3. Click + New Server to open a form that functions as a server creation request.
4. Fill out the required parameters in the form.
5. You will be prompted for permission to create a server access item in your ArcGIS Online account. which will be used by you to manage the dynamic or tiled imagery server. Check the box to confirm and submit the server creation request by clicking Submit Request. This will create a server access item in your ArcGIS Online account.

P&A Access Item

A P&A (or Processing and Analysis) access item is created in ArcGIS Online as an item when an administrator enables a P&A task in ArcGIS Image Dedicated Site.

Complete the following steps to enable a processing and analysis task and create an access item:

1. Verify that you are signed into the ArcGIS Image Dedicated Site and have the necessary privileges to enable or disable P&A tasks.
2. Navigate to the Processing and Analysis tab -> GeoProcessing Task. Expand a folder by clicking on the expandable button to view all available processing and analysis services along with their respective tasks.
3. Choose to enable all tasks within a Processing and Analysis service simultaneously or enable a specific task that is relevant to your needs.
4. Each task is accompanied by a description detailing the type of processing and analysis it performs. Upon enabling a task, you will be prompted to either create a new access item or utilize an existing item for that specific task.
5. It is possible to update the access item of a task at any time by clicking on the "Update Item ID" button.

Designate a Group for ArcGIS Image Dedicated Access

If you need some of your team members to have access, then you can create a new group, share the item within the group, and add team members to that group so they can access it in ArcGIS Online. You can use an existing group if the membership is a good match with your intended ArcGIS Image Dedicated members. If not, you will need to create a new group in ArcGIS Online.

If you have a group with designated members already added to the group, the access items can be shared with the group from the item details page. You can have a single group or multiple groups for all three access items and share the items with these groups, which will provide access to the other members in your organization to administer ArcGIS Image Dedicated, manage dynamic and tiled imagery, and processing and analysis capabilities. If you would like to share an access item with a group, you will need to edit group's sharing abilities for that item.