

AN ESRI
TECHNICAL PAPER

May 2021

ArcGIS Enterprise deployment guide—3D data as scene layers

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



Copyright © 2021 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.

Table of contents

Executive summary 4

Understand your data..... 4

Hardware and software considerations 9

Set up your environment 13

Test your environment 14

Monitor and manage once deployed 14

ArcGIS Enterprise scene layer deployment guide

Executive summary

Scene layers represent large geospatial data in 3D, making it easier for your audience to understand and experience geographic information in a more realistic way. 3D data, though, is usually large and complex. That means you need to understand your data and data storage options, and you need to have the right resources available to store and access this data in an efficient way.

Every use case is individual, and deployment planning needs to be adapted to your specific needs. However, there are recommendations that can help you successfully deploy scene layers. This guide consolidates information about scene layer types and what sort of caching options your publishers should use based on how the scene layers will be used. This guide also provides guidelines on the hardware you may require to successfully store and access scene layers in ArcGIS Enterprise.

The information in this guide is intended for GIS managers responsible for implementing 3D scene layers in an ArcGIS Enterprise deployment and IT staff who are responsible for acquiring, configuring, and maintaining the hardware and software necessary to run scene layers, store scene layer data, and consume scene layers.

Understand your data

Before you configure a deployment for 3D scene layers, review what source data your organization has and the best way to publish it. Estimate how much storage you will need to host the data and decide what scene layer type to use. You will further estimate how many users you are expecting to visit the website and navigate within the scene. All this information is important, as it allows you to include the resources required to make your ArcGIS Enterprise deployment work well with scene layers.

Scene layer types

Publishers can create web scenes in ArcGIS Pro that contain 3D data to be published as a scene layer. Before you configure a deployment for 3D scene layers, decide which layers will participate in web scenes and which layer types will be represented as scene layers.

Scene layers are cached web layers that are optimized for displaying large amounts of 3D features on the web, on mobile devices, and on desktop clients. Scene layers achieve this optimization by showing the content at different levels of detail. Less detail is displayed at larger scales and more detail is displayed as you zoom in to the scene.

Publishers in your organization can create several different types of scene layers depending on their source data and how they want to visualize the data.

The following table describes the type of scene layer you can create for each data source type, and the purpose of each scene layer type. Links in the table take you to pages in the ArcGIS Pro help that provide more information.

Types of data sources	Scene layer types	Example feature types in the scene
Multipatch or 3D object feature class	3D object	3D objects, such as textured or untextured buildings.
Building layer	Building	Complex digital models of buildings, including structural and architectural layers.
Open scene graph binary (OSGB) file or files	Integrated mesh	A triangular interlaced structure representing built and natural 3D features with realistic textures.
Points with z-values (3D points)	Point	Individual 3D points in a layer, to which you can apply a style so the points look like what they represent, such as trees, benches, light poles, and water towers.
LAS, zLAS, LAZ files, or LAS datasets	Point cloud	3D thematic landscapes; the symbology renderer applied to the source data before publishing determines the look of the resultant scene layer.

Table 1: The type of data source you have determines the type of scene layer you publish.

Data complexity

3D data can vary significantly in complexity and, therefore, in performance. The following characteristics can influence the size of the storage needed and the overall performance on ArcGIS clients.

- **Textures**—Textures are the images that are pasted onto the sides of 3D models to give your scene a realistic appearance. 3D object scene layers support textures.

Using textures increases the size of the 3D object scene layer as well as the time it takes to draw in an ArcGIS client. For this reason, textures are often compressed and aggregated in texture atlases to increase performance.

Whether you compress textures depends on the ArcGIS client you want to use. For example, when creating a 3D object scene layer to use in a mobile application, you can set the Texture Optimization option in the [Create 3D Object Scene Layer Package](#) tool with the Texture Optimization parameter set to MOBILE. In this case, the best compression format for mobile use will be applied. Another option is DESKTOP that allows you to add texture compressions optimized for desktop applications. By default, scene layers with textures always include JPEG as format for texture information.

- **Transparency**—Transparency allows you to see features that are behind other features. For example, you can apply transparency to make all windows of a building semitransparent or to make trees partially transparent so you can see the features below and behind the trees.

You can also apply different levels of transparency to represent different values for a feature, similar to how you would use a color ramp. For example, you can apply transparency based on a field to show the percentage of occupation of a building.

In both cases, the scene layer package must store additional information about the transparency and, therefore, applying transparency increases the size of the scene layer.

- **Number of features**—A more obvious reason why the size or performance of a scene layer may change is the number of features it contains. Because the features are organized in levels of detail, the drawing performance when zooming in to the scene layer does not decrease. This makes it possible to show millions of features. How many levels of detail you have and how you define when clients should switch to using the next higher or next lower level of detail also affect drawing performance.
- **Number of vertices in each feature**—Vertices are the invisible points where a line or polygon changes direction. The more vertices a feature contains, the more complex the geometry and, therefore, the more resources are used by clients to draw a feature. For example, if you have a highly detailed model of a building with balconies, details on doors and windows, and columns with elaborate capitals, the geometries will contain more vertices to represent these details correctly.
- **Symbol types**—In a point scene layer, point features are represented by symbols. The complexity of the 3D symbol used for each point increases the resources needed to render the layer. For example, a realistic tree with textures and many individual geometries will need more resources than a simple thematic tree.

Certain symbols provide a different level of detail (LoD) depending on the scale at which the layer is viewed.

The type of symbol you use depends on the type of scene you want to author. Generally, the more realistic the scene is, the more complex the symbols used to represent each point feature are.

- **Physically based rendering (PBR) material**—To achieve the most realistic representation of an object, you can use PBR materials to show such things as the shininess of a car or the roughness of concrete on the façade of a house. You will need more resources to store and draw scene layers using PBR materials.

Usually, PBR materials are used in building scene layers, in 3D object scene layers, or in the symbols representing point scene layers.

- **Number of layers in a web scene**—In ArcGIS, different data is represented in separate layers composing the scene. The more layers that are present in a scene, the more resources are needed to store and render the scene. Similarly, the more complex the layers are, the more resources are needed to store and render the scene. As the number and complexity of layers increase, drawing and query performance decrease.

The following graphic illustrates the relationship between certain data complexity elements and drawing performance for 3D data:

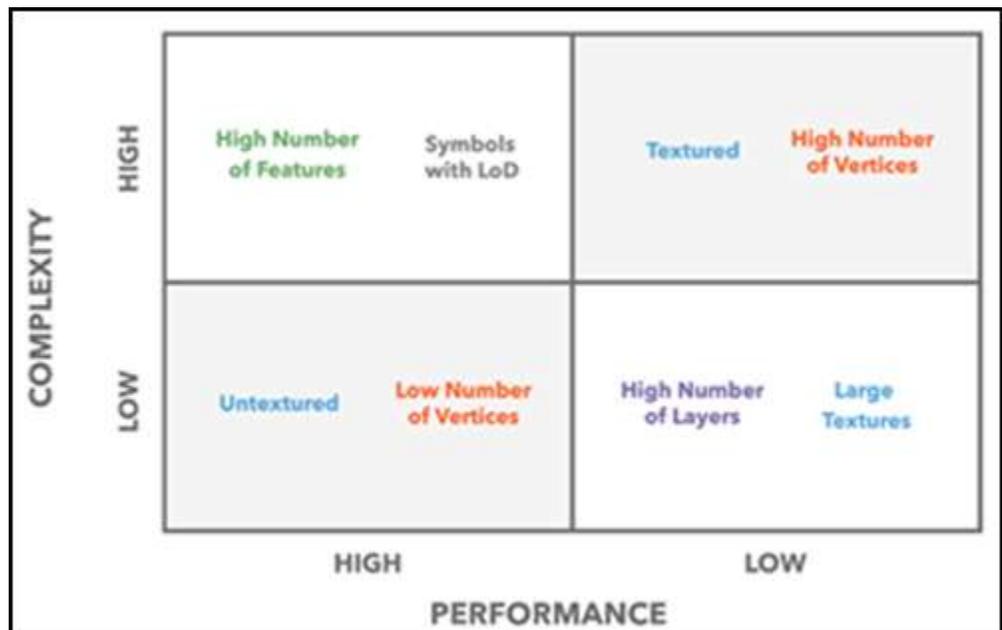


Figure 1: As data complexity increases, drawing performance typically decreases.

Usage patterns

Scene layers contain [cached information](#) organized in levels of detail. This allows ArcGIS clients to rapidly draw scene layers independent of the distance to individual features. How you intend to use the scene layers and maintain their content affects how you cache the data. Therefore, before creating a scene layer, you must decide how it will be used.

Usage pattern options are described in the table on the following page.

Usage pattern	Description	When to use
<p>Static</p>	<p>You can use this type of scene layer to visualize 3D content by changing symbology, identify features and view the attribute information, or exclude features based on a selection.</p> <p>You cannot edit the content of the scene, but the layer can be replaced by a newer version of a scene layer.</p> <p>To create a static scene layer, you create a scene layer package (SLPK) or cache the scene layer locally on the client when publishing.</p>	<p>This pattern is often used if data is shared outside of an organization or access to the source data needs to be restricted.</p> <p>This pattern also allows you to share a subset of features with your audience. For example, you can define queries for multipatch or point feature layers or clip an integrated mesh scene layer to create a scene layer package.</p>
<p>Dynamic</p>	<p>Feature-based scene layers—such as 3D object scene layers, building scene layers, and point scene layers—can be associated with feature data stored in the system of record within your GIS system. The associated feature data can be accessed through the feature service, or feature data can be copied to the relational data store and accessed via a hosted feature layer.</p> <p>A scene layer and its associated feature layer build a unit that allows you to use the capabilities of the feature layer through the scene layer.</p> <p>This requires you to cache the scene on the server when publishing.</p>	<p>If you need to update individual feature geometries or attributes, your usage pattern is dynamic.</p> <p>Besides allowing editing on the scene layer, you may also want to use a scene layer with an associated feature layer if you need to dynamically query the data. For example, the ArcGIS Enterprise Scene Viewer uses the associated feature layer to search within a scene layer.</p>

Table 2: If 3D data is relatively static, you can package the data or cache scene data locally when publishing; if data is dynamic, cache on the server when publishing and create an associated feature layer.

Caching options

Scene layers use cached data to provide better drawing performance. Those who publish scene layers must choose where scene caches are built. This choice should be based on the usage pattern for the layer, as discussed in the previous section.

When publishing from ArcGIS Pro, the publisher has the following two options for storing the scene layer cache: **Cache on the server** or **Cache locally**.

- Choose the **Cache on the server** option if you publish a 3D object, building, or point scene layer and expect to update the content on a regular basis.

When the publisher chooses the **Cache on the server** option, a scene layer and an associated feature layer are created in the ArcGIS Enterprise organization. The feature layer contains an editable vector representation of the data present in the scene layer. Having an associated feature layer allows edits made to the feature layer to be used to update the scene cache.

The publisher has the additional option to leave the feature layer data in the registered data source (reference registered data) or copy the data to a hosted feature layer. To leave the feature layer data in a registered data source, the data source must be an enterprise geodatabase.

- Choose the **Cache locally** option if the scene layer does not need to be updated frequently. When the publisher chooses this option, no associated feature layer is created. An SLPK and scene layer are created in your ArcGIS Enterprise organization.

This is the only option available for integrated mesh and point cloud scene layers.

If the contents of the scene layer cache do need to be updated, the layer owner can replace the old scene layer with a new scene layer.

The presence or absence of an associated feature layer, and whether the feature data is copied or stays in place, impacts storage, network traffic, and the time it takes to create a scene layer. With the same number of features and feature complexity, the amount of time it takes to publish increases in the order shown below.

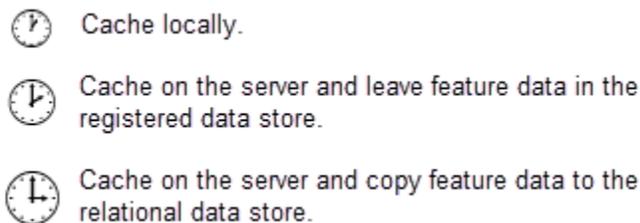


Figure 2: Caching locally takes less time; caching on the server and copying data takes the most time.

No matter which option the publisher chooses when publishing from ArcGIS Pro, the scene layer cache is stored in the tile cache data store in the ArcGIS Enterprise deployment.

Hardware and software considerations

When planning for an ArcGIS Enterprise deployment to support hosted scene layers, you need to take into consideration the resources you need—such as software, licensing, hardware (physical or virtual), and hardware capacity—and logistics—such as deployment options and privileges.

General guidelines

Hardware considerations are difficult to define without knowing the exact scenario for which the architecture is planned. Often this environment needs to accommodate many different scenarios. Definitive statements cannot be made about the size or

number of machines in your deployment or the options you use, because each implementation is different. Therefore, general guidelines for these pieces of an implementation are provided. You can review the minimum hardware requirements for [ArcGIS Server](#), [Portal for ArcGIS](#), and [ArcGIS Data Store](#), but these are the minimum requirements to install and start these components; they do not take into consideration the number of services, items, and data you will create. To help determine actual resources needed, complete the baseline tests in the technical paper [ArcGIS Enterprise deployment guide—Scene layer benchmark testing](#).

Tip: *If you need assistance configuring a base ArcGIS Enterprise deployment, contact Esri consulting services (in the US) or contact your Esri distributor (outside the US) to get an in-depth analysis of your hardware needs.*

The software and privileges required to publish a scene layer are defined in the following sections.

Software

To publish hosted scene layers, you must have the following software installed and configured:

- [ArcGIS Pro](#)

- [A base ArcGIS Enterprise deployment](#)

A base deployment includes Portal for ArcGIS, an ArcGIS GIS Server site (licensed as standard or advanced) configured as the hosting server, two installations of ArcGIS Web Adaptor, and ArcGIS Data Store with relational and tile cache data stores configured.

The tile cache data store is critical for 3D workflows in ArcGIS Enterprise. In most production environments, you will install each of these components on separate machines or sets of machines if you will be configuring a highly available ArcGIS Enterprise deployment.

- An [enterprise geodatabase](#) and [federated ArcGIS GIS Server](#) (optional)
If you will publish from ArcGIS Pro with the option to **Cache on the server** and feature data will remain in a registered database, you need an enterprise geodatabase in which to store the feature data. You'll register the geodatabase with the ArcGIS GIS Server site before publishing. When you publish, an ArcGIS Server feature service is created with the scene layer.

To run hosted layers and ArcGIS Server web services on separate sites, you need a federated ArcGIS GIS Server. If you upload an existing SLPK to the portal and publish, you do not require a registered database.

Tip: *When installing software components that need to communicate with each other, the best practice is to install the software in the same data center. If you deploy ArcGIS Enterprise in the cloud, all virtual machines need to be in the same cloud region. This reduces the network traffic between the components and ensures better performance.*

Publisher privileges

To publish a hosted scene layer, the organization member must be assigned a Creator or higher level [user type](#) and either be a member of the [default Publisher role](#) or a member of a [custom role](#) that has privileges to create content and publish hosted scene layers.

In addition, if publishing from ArcGIS Pro with the option to **Cache on the server**, the custom role requires a privilege to publish a feature layer. If the feature data will be copied when publishing, the custom role requires the privilege to publish hosted feature layers. If feature data will remain in the registered database when publishing, the custom role requires the privilege to **publish server-based layers**.

Storage space

The amount of storage space needed and where it is needed varies depending on the [caching options](#) used when publishing. How you configure tile cache [data store backups](#) also affects storage space and locations.

Some general guidelines about the minimum amount of storage space you need for scene layers published from ArcGIS Pro with the option to **Cache locally** (which creates an SLPK behind the scenes) or when you upload an SLPK file to the portal website and publish are as follows:

- Temporary files are created on the hosting server and ArcGIS Enterprise portal machines that require there to be free space on the hosting server machines amounting to approximately the same size as the SLPK package and free space on the ArcGIS Enterprise portal machines equivalent to two times the SLPK package size.
- The amount of space needed in the tile cache data store is equivalent to two times the size of the SLPK file you publish.

Note: *Once you publish, you can delete the SLPK file from your portal if you do not need to keep it. (For example, if others in your organization need to download and use the file, don't delete it.) Deleting the SLPK file from the portal saves storage space on your portal machines.*

When you publish with the option to **Cache on the server** and copy feature data, the following approximate minimum storage values apply:

- You need five times the size of the source data available in free space on the hosting server machines.
- A large amount of temporary space is required when creating scene caches on the server. If running the hosting server on Windows, the default temp folder on your C: drive is used. If space on your C: drive is limited, specify an alternate temp location that contains enough space. If running the hosting server on Linux, be sure you configure a large enough temp directory. See [Configure directories for local cache generation](#) in the ArcGIS Enterprise help for more information on setting a temp file location for cache creation.
- On the ArcGIS Pro client machine, you need two times the size of the source data for temporary data storage.

- The number of features and feature complexity will affect the amount of storage space needed in the relational data store for the feature data. Point data takes the least space, simple lines and polygons take more space than points, and complex lines and polygons (those with many vertices) take the most space.
- The amount of space needed in the tile cache data store varies depending on the complexity of the data and whether textures are present. Simple data with no textures can use the size of the cache + the same amount of space for temporary storage. To help you determine the overall cache size for your data, publish a smaller subset of the data and extrapolate the cache size for the entire dataset from that sample cache. Tests run by Esri found that, for a multipatch feature class with textures, the cache size is approximately five times the size of the feature class.

The following table summarizes the approximate minimum free disk space required per machine:

	Primary portal machine	Hosting server machines	Tile cache data store machines	ArcGIS Pro machine
Publish from a scene layer package (SLPK) or cache locally.	The size of the SLPK to store it as an item in the portal + the size of the SLPK in temp space = 2x the size of the SLPK.	The size of the SLPK in temp space.	The size of the SLPK for the tile cache + the same amount of storage in temp space = 2x the size of the SLPK.	When caching on the client, 2 times the size of the SLPK in temp space.
Publish from a geodatabase and cache on the server.*	Not applicable.	2x the size of the source data + 3x the size of the source data in temp space = 5x the size of the source data.	The cache size + the same amount of storage in temp space. Cache size can be up to 5x the size of the source data.	The size of the source data.

Table 3: Free disk space is required on the portal, hosting server, tile cache data store, and ArcGIS Pro machines.

*If feature data remains in a registered database, no additional space is needed for features. If the feature data is copied to the relational data store, account for storage space needed in the relational data store.

Memory and CPU

Memory and CPU values need to be higher on the machines where you will be storing and working with 3D data because of the data's larger size and complexity.

By investing in your enterprise hardware, you make it possible for client applications to more easily and quickly access the 3D data.

ArcGIS system requirements pages list the minimum resources necessary to install and start the software. To facilitate high performance of scene layers, use hardware that possesses resources that exceed these minimum system requirements. Depending on your scenario, you may need different requirements for different ArcGIS Enterprise components.

For example, the hosting server and tile cache data store machines require the most memory and CPU. The machines you use for these software components should have a minimum of 32 GB of RAM.

To decrease the amount of time it takes to create the scene layer caches when caching on the server, a minimum of 8 cores on the hosting server is recommended. When clients consume the scene layers, you need 8 cores on the tile cache data store machines.

Set up your environment

After choosing your hardware and software configuration, you need to install the different ArcGIS components. Have your IT department make the machines available that you require. Next, you or the IT department staff can install the software.

Install software

See the following documentation for instructions on installing ArcGIS software:

- The Set up [ArcGIS Pro](#) (minimum 2.6) section of the ArcGIS Pro help.
- Installation guides for [Portal for ArcGIS](#), [ArcGIS Server](#), [ArcGIS Web Adaptor](#), and [ArcGIS Data Store](#) (minimum 10.8.1). Use the **Other versions** list to access the guide specific to your operating system.

Configure data store

Scene layer caches are stored in the tile cache data store. Therefore, your ArcGIS Enterprise deployment must contain a tile cache data store to allow members of your ArcGIS Enterprise organization to publish hosted scene layers. If your deployment doesn't currently include a tile cache data store, install ArcGIS Data Store on at least one additional machine and configure a tile cache data store. For best performance, install the tile cache data store on a machine that does not have another data store or database running on it; this avoids having to share resources with other components in the deployment, which can degrade performance or even lead to the tile cache data store shutting down due to lack of free disk space.

Tip: *If any of the machines where you install ArcGIS Data Store contain more than one network interface controller (NIC) card or multiple DNS entries, configure the ArcGIS Data Store installations to identify which IP address and host name to use for network communication. See the instructions in [Create a data store](#).*

The number of machines in your tile cache data store and the mode you run it in varies by your organization's needs. If you require high availability, deploy the tile cache data store in primary-standby mode, in which case, the tile cache data store will contain two machines. If you anticipate that you will need to add machines to the tile cache data store to accommodate large amounts of scene layer caches, run the tile cache data store in cluster mode. See the next section and [Tile cache data store deployment modes](#) in the ArcGIS Enterprise help for more information.

Alter default settings

By default, tile cache data stores run in primary-standby mode. You might start with a single or two-machine tile cache data store running in primary-standby mode but later find your publishers are creating a lot of scene layers. In that case, you can change to a cluster mode and add more machines to the data store.

1. To switch to cluster mode, run the [configuredatastore utility](#) with the --mode operation set to cluster.
2. Set up new machines and [configure them to be part of the existing tile cache data store](#).
3. If necessary, [rebalance caches for existing scene layers](#). If you do not rebalance existing caches, new scene layers will use the new machines, but existing layers will remain on the machine or machines that were in your tile cache data store before you added machines to it.

No backups are created of tile cache data stores until you configure them. You should [configure a backup location](#) for the tile cache data store before people start publishing scene layers.

Test your environment

Once you understand your data, have estimated the resources you will need, and have configured the required software, test your environment. See the technical paper, [ArcGIS Enterprise deployment guide—Scene layer benchmark testing](#), for sample test scenarios and data you can use to establish a performance baseline for your system.

Monitor and manage once deployed

After you create your deployment, you must monitor it to be sure it continues to perform as expected. Even if initial performance testing of your environment shows that your hardware and software configuration is adequate, usually more and more web services are published over time. You need to keep an eye on the system to ensure disk space and memory are not running low and no failures have occurred. You also need to implement a backup schedule and data recovery plan, and test that your plan works.

Tip: For information on how to manage issues you may encounter when configuring or maintaining the tile cache data store, see [Troubleshoot ArcGIS Data Store](#) in the ArcGIS Enterprise help. For information on how to manage issues that users may encounter when publishing or using scene layers, see [Troubleshoot scenes and 3D data](#).

Backups and recovery

Backups and a tested recovery plan are important for all production systems.

When you deploy ArcGIS Enterprise on-premises, configure backups to write to a file share in the in the same data center. If you create backups to a location outside of where the tile cache data store machines are stored, the backup process will take longer. If you want to copy backups to a location outside your data center to guard against catastrophes such as destruction of the data center, automate the copying of files from the file share to place them in another location, such as in cloud storage.

If you deploy ArcGIS Enterprise in a cloud environment, configure backups to write to a cloud storage location in the same region as ArcGIS Enterprise.

You can use the `webgisdr` utility installed with Portal for ArcGIS to create backups of most components in your deployment. See [ArcGIS Enterprise backups](#) in the ArcGIS Enterprise help for more information on using this utility.

ArcGIS Data Store backups

You should back up the entire ArcGIS Enterprise deployment, but you can also configure the tile cache data store to automatically generate backup files on a schedule that you specify. These backup files should be stored on machines separate from the data store and other components in the ArcGIS Enterprise deployment. See [Manage data store backups](#) in the ArcGIS Enterprise help for information on how to configure backup schedules and locations.

Monitor machine health

The scene cache data, feature data (if the scene has an associated feature layer), services, and associated items in the Enterprise portal must all be available for the scene layer to be available. Therefore, you need to ensure that all machines in the ArcGIS Enterprise deployment continue to function.

ArcGIS Data Store

You can [monitor ArcGIS Data Store log files](#) for information on data store health. Logs are written to the `<ArcGIS Data Store directory>\logs\<machine name>` directory. The default location is `C:\arcgisdatastore\logs\<machine name>` on Microsoft Windows machines and `datastore/logs/<machine name>` on Linux machines. The log file specific to tile cache data stores is named `couchlog`.

If the scene layers you publish have associated hosted feature layers, monitor the logs for the relational data store machines as well. The log file specific to relational data stores is named `database`.

If the scene layers you publish have associated feature layers with data in a registered database, use the tools available with the database management system where the data is stored to monitor database health.

If the data store logs include messages indicating a machine is running low on disk space, you need to add more space to the existing machines or add machines to the

tile cache data store. Adding machines may require you to change the tile cache data mode.

When a relational or tile cache data store machine contains less than 10 GB of free disk space, ArcGIS Data Store begins logging warnings that you are running out of disk space. Once the disk drive contains less than 1 GB of free space, the tile cache data store is shut down. When the tile cache data store shuts down, no one can access existing layers or publish new ones.

ArcGIS Server

You can access ArcGIS Server logs from ArcGIS Server Manager. Sign in to ArcGIS Server Manager for your hosting server site and click **Logs > View Logs**. When free disk space on the drive where the ArcGIS Server installation directory or log directory resides falls below 10 GB, messages to that effect are written to the log file.

For more information about ArcGIS Server logs, including how to alter the logging level and delete logs, see [Work with server logs](#) in the ArcGIS Enterprise help.

Portal for ArcGIS

Most troubleshooting information will appear in the ArcGIS Server and ArcGIS Data Store logs. However, if publishing fails because the SLPK file could not be uploaded, or if the publishing operation takes so long that the portal token expires, those messages will appear in the Portal for ArcGIS logs. See [Work with portal logs](#) and related help pages in the ArcGIS Enterprise help for information about accessing logs and altering logging levels.

About the authors

This paper was composed by staff from multiple teams across Esri:

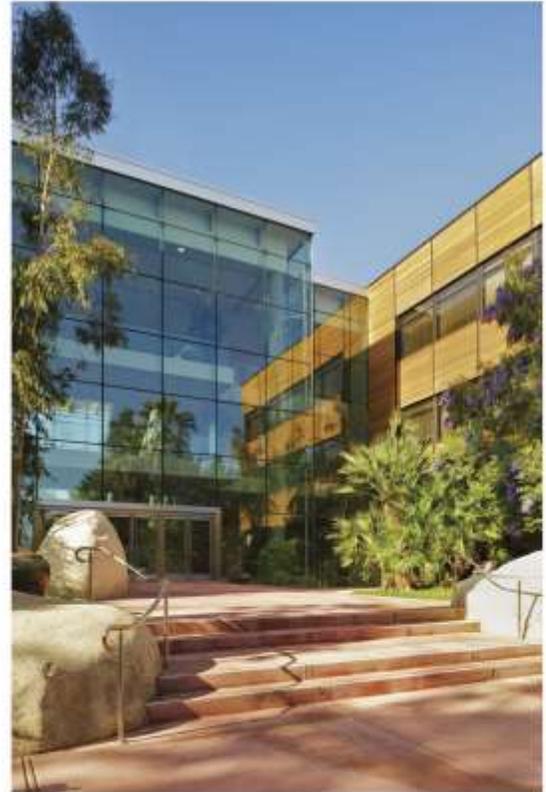
- Aaron Lopez, Advanced Enterprise Systems
- Ivonne Seler, PE, ArcGIS Pro
- Jill Edstrom-Shoemaker, PM, ArcGIS Enterprise
- Kimberly Peter, PE, ArcGIS Enterprise
- Philip Heede, PM, ArcGIS Enterprise
- Ricardo Perez, Technical Consultant, Advanced Enterprise Systems
- Richard Vargas, PE, ArcGIS Pro
- Satish Malipeddi, PE, ArcGIS Enterprise
- Sean Morrish, PE, ArcGIS Pro
- Zoe Veale, PE, ArcGIS Pro



Esri, the global market leader in geographic information system (GIS) software, offers the most powerful mapping and spatial analytics technology available.

Since 1969, Esri has helped customers unlock the full potential of data to improve operational and business results. Today, Esri software is deployed in more than 350,000 organizations including the world's largest cities, most national governments, 75 percent of Fortune 500 companies, and more than 7,000 colleges and universities. Esri engineers the most advanced solutions for digital transformation, the Internet of Things (IoT), and location analytics to inform the most authoritative maps in the world.

Visit us at esri.com.



Contact Esri

380 New York Street
Redlands, California 92373-8100 USA

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

Offices worldwide
esri.com/locations

For more information, visit
esri.com/URL.