

ArcGIS® Enterprise: Administration Workflows



ArcGIS[®] Enterprise: Administration Workflows

STUDENT EDITION

Copyright © 2019–2025 Esri

All rights reserved.

Course version 7.0. Version release date December 2025.

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: Director, Contracts and Legal, Esri, 380 New York Street, Redlands, CA 92373-8100, USA.

Export Notice: Use of these Materials is subject to U.S. export control laws and regulations including the U.S. Department of Commerce Export Administration Regulations (EAR). Diversion of these Materials contrary to U.S. law is prohibited.

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

Commercial Training Course Agreement Terms: The Training Course and any software, documentation, course materials or data delivered with the Training Course is subject to the terms of the Master Agreement for Products and Services found at the following website: esri.com/TrainingTerms. The license rights in the Master Agreement strictly govern Licensee's use, reproduction, or disclosure of the software, documentation, course materials and data. Training Course students may use the course materials for their personal use and may not copy or redistribute for any purpose. Contractor/Manufacturer is Esri, 380 New York Street, Redlands, CA 92373-8100, USA.

Esri Marks: Esri marks and product names mentioned herein are subject to the terms of use found at the following website: esri.com/EsriMarks.

Other companies and products or services mentioned herein may be trademarks, service marks, or registered marks of their respective mark owners.

Table of Contents

| | |
|--|------|
| Esri resources for your organization | viii |
|--|------|

Course introduction

| | |
|-----------------------------------|---|
| Course introduction | 1 |
| Course goals | 1 |
| Installing the course data | 1 |
| Icons used in this workbook | 3 |

1 Business needs and ArcGIS Enterprise solutions

| | |
|---|------|
| Lesson introduction | 1-1 |
| Mapping business needs to a system deployment..... | 1-2 |
| Architecting an ArcGIS system..... | 1-3 |
| Explore the ArcGIS system patterns | 1-5 |
| Providing capabilities to support system patterns | 1-7 |
| Using the ArcGIS Enterprise Functionality Matrix..... | 1-9 |
| Business needs and ArcGIS Server licensing roles..... | 1-11 |
| Lesson review | 1-14 |
| Answers to lesson 1 questions..... | 1-15 |

2 Adding capabilities through federation

| | |
|---|------|
| Lesson introduction | 2-1 |
| Base ArcGIS Enterprise deployment..... | 2-2 |
| Federating additional server sites | 2-4 |
| Explore considerations for federating additional server sites | 2-8 |
| Exercise environment..... | 2-10 |
| Exercise 2: Federate an Image Server site..... | 2-12 |
| Examine a stand-alone Image Server site..... | 2-13 |
| Investigate a raster store..... | 2-14 |
| Federate the Image Server site | 2-15 |
| Examine the effects of federation | 2-17 |
| Use the raster analysis tools | 2-19 |
| Lesson review | 2-23 |

Answers to lesson 2 questions..... 2-24

3 Ensuring resilient scaling through distributed computing

Lesson introduction 3-1

Distributed computing and business needs 3-2

Explore ArcGIS Enterprise high availability 3-4

Making ArcGIS Enterprise components highly available 3-6

Explore workload separation 3-13

Comparing high availability and workload separation 3-14

Workload separation by server capability 3-15

Exercise 3: Implement workload separation 3-17

 Configure a raster store for second Image Server site 3-18

 Federate the Image Server site 3-20

 Explore system raster tools..... 3-22

 Use Image Server to perform distributed raster analysis 3-24

 Inspect running raster analysis services 3-28

 Explore imagery layer output..... 3-30

Observing a complex ArcGIS system 3-32

Lesson review 3-35

Answers to lesson 3 questions 3-36

4 Web service fundamentals

Lesson introduction 4-1

Web service fundamentals 4-2

Explore user-managed data and ArcGIS-managed data 4-4

How layers enable efficient exposure of services 4-6

Understand the relationship between services and layers 4-7

Exercise 4: Compare layers, services, and data 4-9

 Publish a feature layer that uses ArcGIS-managed data 4-10

 Investigate a feature layer that uses ArcGIS-managed data 4-12

 Publish a feature layer with user-managed data 4-14

 Investigate a feature layer that uses user-managed data 4-16

Lesson review 4-18

Answers to lesson 4 questions..... 4-19

5 User-managed data

Lesson introduction 5-1
Types of user-managed data 5-2
Referencing registered data from ArcGIS Pro 5-5
Capabilities for services that use user-managed data 5-7
Versioning 5-8
Exercise 5: Perform branch versioning workflows 5-11
 Register a dataset as branch versioned 5-12
 Share a branch-versioned layer 5-14
 Inspect branch management capabilities in ArcGIS Server Manager 5-17
 Create versions 5-19
Lesson review 5-22
Answers to lesson 5 questions..... 5-23

6 Optimizing services that reference user-managed data

Lesson introduction 6-1
Service instances..... 6-2
Service instance types..... 6-8
Exploring service instances 6-10
Consider when to use shared instances 6-11
Exercise 6: Work with shared instances 6-13
 Inspect ArcSOC processes using Task Manager 6-14
 Change shared instance settings for a server site 6-16
 Configure a service to use shared instances..... 6-17
 Explore deletion behavior of a non-hosted feature layer..... 6-19
Reviewing ArcGIS Server logs and statistics reports 6-23
Troubleshooting common service performance issues..... 6-26
Lesson review 6-28
Answers to lesson 6 questions..... 6-29

7 ArcGIS-managed data

Lesson introduction 7-1
ArcGIS Data Store 7-2

| | |
|--|------|
| Types of ArcGIS-managed data | 7-4 |
| Relationship between data storage options and service type | 7-6 |
| ArcGIS Data Store backup options | 7-8 |
| ArcGIS Data Store backup defaults | 7-9 |
| Managing ArcGIS Data Store using command utilities | 7-11 |
| Exercise 7: Manage ArcGIS Data Store backups | 7-14 |
| Create a shared network folder | 7-15 |
| Run a command utility | 7-16 |
| Configure relational data store backup properties | 7-18 |
| Take a manual backup | 7-20 |
| Lesson review | 7-21 |
| Answers to lesson 7 questions | 7-22 |

8 Optimizing services that use ArcGIS-managed data

| | |
|--|------|
| Lesson introduction | 8-1 |
| Services that use or interact with ArcGIS-managed data | 8-2 |
| Optimizing hosted feature services using the object store | 8-5 |
| Distributed data enables distributed analysis | 8-8 |
| Exploring distributed raster analytics results | 8-10 |
| Determining a service optimization strategy | 8-11 |
| Exercise 8: Use ArcGIS Server to optimize feature analysis | 8-13 |
| Inspect SpatialAnalysisTools service instances | 8-14 |
| Increase instances for the SpatialAnalysisTools service | 8-15 |
| Set up the analysis | 8-17 |
| Run the analysis | 8-18 |
| Inspect hosted feature layer output | 8-19 |
| Explore deletion behavior of a hosted feature layer | 8-20 |
| Lesson review | 8-21 |
| Answers to lesson 8 questions | 8-22 |

9 The spatiotemporal big data store

| | |
|---|-----|
| Lesson introduction | 9-1 |
| Understanding the spatiotemporal big data store | 9-2 |
| ArcGIS Enterprise and the spatiotemporal big data store | 9-5 |
| How the spatiotemporal big data store relates to other ArcGIS Enterprise components | 9-8 |

| | |
|---|------|
| Exercise 9: Explore data in the spatiotemporal big data store | 9-9 |
| Investigate the spatiotemporal big data store | 9-10 |
| View a spatiotemporal big data store node | 9-11 |
| Explore the spatiotemporal big data store in GeoEvent Manager | 9-12 |
| View data stored in the spatiotemporal big data store..... | 9-14 |
| Manipulate data in a spatiotemporal big data store | 9-18 |
| Distributed data and the spatiotemporal big data store | 9-21 |
| Distributed analysis and the spatiotemporal big data store | 9-24 |
| Lesson review | 9-26 |
| Answers to lesson 9 questions..... | 9-27 |

10 Configuring a real-time analysis solution

| | |
|--|-------|
| Lesson introduction | 10-1 |
| Real-time data | 10-2 |
| Working with real-time data | 10-3 |
| What is a GeoEvent? | 10-4 |
| Explore configuring real-time analytics | 10-6 |
| Building a GeoEvent Service in GeoEvent Manager..... | 10-7 |
| Exercise 10: Configure a GeoEvent Service to write real-time data to a spatiotemporal big data store | 10-8 |
| Federate a GeoEvent Server site | 10-9 |
| Investigate an existing input connector..... | 10-10 |
| Configure the GeoEvent Simulator | 10-11 |
| Investigate GeoEvent Definitions | 10-14 |
| Create a new output connector | 10-15 |
| Create a GeoEvent Service..... | 10-17 |
| Observe data as it flows through a GeoEvent Service..... | 10-21 |
| Lesson review | 10-26 |
| Answers to lesson 10 questions..... | 10-27 |

11 Distributed collaboration

| | |
|---------------------------------------|------|
| Lesson introduction | 11-1 |
| Sharing content between portals | 11-2 |
| Distributed collaboration..... | 11-3 |
| Planning for a collaboration | 11-5 |

| | |
|---|-------|
| Steps to create a distributed collaboration | 11-7 |
| Explore sharing content in a collaboration | 11-9 |
| Exercise 11: Use a collaboration to share content | 11-10 |
| Ensure HTTPS only access to each portal | 11-11 |
| Review the trusted certificates | 11-12 |
| Create the collaboration invitation | 11-14 |
| Accept the collaboration invitation | 11-17 |
| Join a workspace from a guest portal | 11-19 |
| Create a new workspace..... | 11-20 |
| Lesson review | 11-24 |
| Answers to lesson 11 questions..... | 11-25 |

12 Automating ArcGIS Enterprise

| | |
|--|-------|
| Lesson introduction | 12-1 |
| Automating administrative tasks..... | 12-2 |
| Automation strategies..... | 12-3 |
| ArcGIS Notebooks and ArcGIS Notebook Server | 12-4 |
| Exercise 12A: Create webhooks as an automation strategy in ArcGIS Enterprise | 12-6 |
| Federate an ArcGIS Notebook Server site..... | 12-7 |
| Create a custom role for notebook authors | 12-8 |
| Upload a notebook | 12-10 |
| Create a webhook | 12-11 |
| Test the webhook..... | 12-13 |
| Working with the ArcGIS API for Python | 12-14 |
| Architecture considerations for ArcGIS Notebook Server..... | 12-16 |
| Exercise 12B: Use ArcGIS Notebooks to schedule administrative tasks in ArcGIS Enterprise | 12-18 |
| Import a notebook..... | 12-19 |
| Run a notebook..... | 12-20 |
| Schedule a notebook task | 12-23 |
| Investigate a user workspace | 12-24 |
| Inspect notebook containers | 12-25 |
| Lesson review | 12-27 |

Answers to lesson 12 questions..... 12-28

Appendices

Appendix A: Data license agreement.....A-1
Appendix B: Answers to lesson review questions..... B-1
Appendix C: Additional resources C-1
Appendix D: Lesson stories.....D-0

Esri resources

Take advantage of these resources to develop ArcGIS® software skills, discover applications of geospatial technology, and tap into the experience and knowledge of the ArcGIS community.

Instructor-led and e-Learning resources

Esri® instructor-led courses and e-Learning resources help you develop and apply ArcGIS skills, recommended workflows, and best practices. View all training options at esri.com/training/catalog/search.

Planning for organizations

Esri training consultants partner with organizations to provide course recommendations for job roles, short-term training plans, and workforce development plans. Contact an Esri training consultant at training@esri.com.

Esri technical certification

The Esri Technical Certification Program recognizes qualified individuals who are proficient in best practices for using Esri software. Esri exams are designed to validate expertise with ArcGIS products, such as ArcGIS Online, ArcGIS Pro and extensions, ArcGIS Enterprise and extensions, and developer technologies. Certifications are offered at three levels: foundation, associate, and professional. Learn more at esri.com/training/certification.

Social media and publications

X (formerly Twitter): [@EsriTraining](https://twitter.com/EsriTraining) and [@Esri](https://twitter.com/Esri)

Esri on LinkedIn: [linkedin.com/company/esri](https://www.linkedin.com/company/esri)

Esri training blog: esri.com/trainingblog

Esri publications: Access online editions of ArcNews, ArcUser, and ArcWatch at esri.com/esri-news/publications

Esri training newsletter: Subscribe at go.esri.com/training-news

Other Esri newsletters: Subscribe to industry-specific newsletters at go.esri.com/subscribe

Esri resources (continued)

Esri Press

Esri Press publishes books on the science and technology of GIS in numerous public and private sectors. esripress.esri.com

GIS bibliography

A comprehensive index of journals, conference proceedings, books, and reports related to GIS, including references and full-text materials. gis.library.esri.com

ArcGIS documentation and tutorials

In-depth information, tutorials, and documentation for ArcGIS products.

ArcGIS Online: arcgis.com

ArcGIS Desktop: desktop.arcgis.com

ArcGIS Enterprise: enterprise.arcgis.com

ArcGIS Pro: pro.arcgis.com

Esri Community

Join the online community of GIS users and experts. community.esri.com

Esri events

Esri conferences and user group meetings offer a great way to network and learn how to achieve results with ArcGIS. esri.com/events

Esri videos

View an extensive collection of videos by Esri leaders, event keynote speakers, and product experts. mediaspace.esri.com

Esri resources (continued)

ArcGIS for Personal Use

Improve your GIS skills at home and use ArcGIS to enhance your personal projects. The ArcGIS for Personal Use program includes a 12-month term license for ArcGIS Pro, extension products, and an ArcGIS Online named user account with 100 service credits. [esri.com/personaluse](https://www.esri.com/personaluse)

GIS Dictionary

This term browser defines and describes thousands of GIS terms. support.esri.com/en-us/gis-dictionary

Course introduction

This course will provide the fundamentals for installing and configuring ArcGIS Enterprise components, including many of the server licensing roles beyond the roles of the GIS (hosting) server. Accepted best practices are presented for all facets of the system, and administrative decisions are based on expressed business needs.

Beyond extending the base ArcGIS Enterprise deployment architecture to encompass additional server licensing roles, considerations are presented for optimizing and maintaining performance and extending end-user capabilities. ArcGIS Data Store is explored in detail, covering storage types and options, as well as fault tolerance and other capabilities.

Various administrator workflows are also discussed, including managing portal members, establishing trust for distributed collaborations, web service tuning, and so on. Finally, methods and opportunities for automation are presented to make administrators more efficient, productive, and with more repeatable results.

Course goals

After completing this course, you will be able to perform the following tasks:

- Apply best practices to configure GIS resources and services.
- Maintain system performance using workload separation and other best practices.
- Configure distributed collaboration between multiple ArcGIS Enterprise portals.
- Use ArcGIS Notebooks and the ArcGIS API for Python to automate common administrative functions.

Installing the course data

Some exercises in this workbook require data. Depending on the course format, the data is available on a DVD in the back of a printed workbook or as a data download. To use the data, extract it to your C:\EsriTraining folder.



DISCLAIMER: Some courses use sample scripts or applications that are supplied either on the DVD or on the Internet. These samples are provided "AS IS," without warranty of any kind, either express or implied, including but not limited to, the implied warranties of merchantability, fitness for a particular purpose, or noninfringement. Esri and its authorized distributors shall not be liable for any damages under any theory of law related to the licensee's use of these samples, even if Esri and its authorized distributors are advised of the possibility of such damage.

Icons used in this workbook



Notes indicate additional information, exceptions, or special circumstances about specific course topics.



Recommended practices improve efficiency and save time.



Esri Academy resources provide more in-depth training on related topics.



Additional resources provide additional information about related topics.



Warnings indicate potential problems or actions that should be avoided.

1

Business needs and ArcGIS Enterprise solutions

Modern GIS workflows are supported by your system architecture. In this lesson, you will explore some architectural best practices to understand how they can provide the framework for understanding which GIS capabilities are needed to support your organization's business requirements. After you have delineated your needed capabilities, you can begin extending the functionality and capacity of your base ArcGIS Enterprise deployment. Sometimes, that may mean scaling out the deployment with additional GIS Servers. In other cases, that may mean adding additional ArcGIS Server licensing roles and server extensions to support the desired capabilities.

Understanding the capabilities that you need is not a simple task. Organizations should conduct a thorough needs assessment to identify which business functions could benefit from expanded GIS capabilities. These benefits could be as simple as augmenting Microsoft Office applications with maps or as extensive as analyzing extremely large datasets for spatial relationships that can provide actionable information.

Topics covered

Aligning architectural best practices to meet business needs

ArcGIS Server licensing roles

Mapping business needs to a system deployment

Creating and maintaining an ArcGIS Enterprise deployment that optimally supports an organization's business goals requires planning and execution. Such an endeavor connects users performing business operations to ArcGIS software that enables analytics, data management, and self-service mapping.

The goal is to map an organization's business goals and operations to a fully functional ArcGIS Enterprise deployment that can support existing business needs. ArcGIS Enterprise is a rapidly evolving system that is continually integrating the latest technologies to help solve GIS problems. As a result, ArcGIS Enterprise provides organizations with the tools to address future business needs as they arise.



Figure 1.1. ArcGIS Enterprise enables organizations to meet their business needs, rather than adapting their workflows to GIS technology.

Provide an example of how your organization translated a business need into a technical implementation.

Architecting an ArcGIS system

An ArcGIS Enterprise deployment provides an ArcGIS system for your organization. Before architecting an ArcGIS system, it is important to consider the capabilities that are needed to support your organization's business needs. These capabilities are provided through the eight ArcGIS system patterns, which are common types of systems that are built with ArcGIS.

Architecture of the ArcGIS system

After the required ArcGIS system patterns are identified, the architecture of the ArcGIS system can be designed to implement those system patterns. The overall architecture of an ArcGIS system consists of three tiers: applications (apps), services, and data.

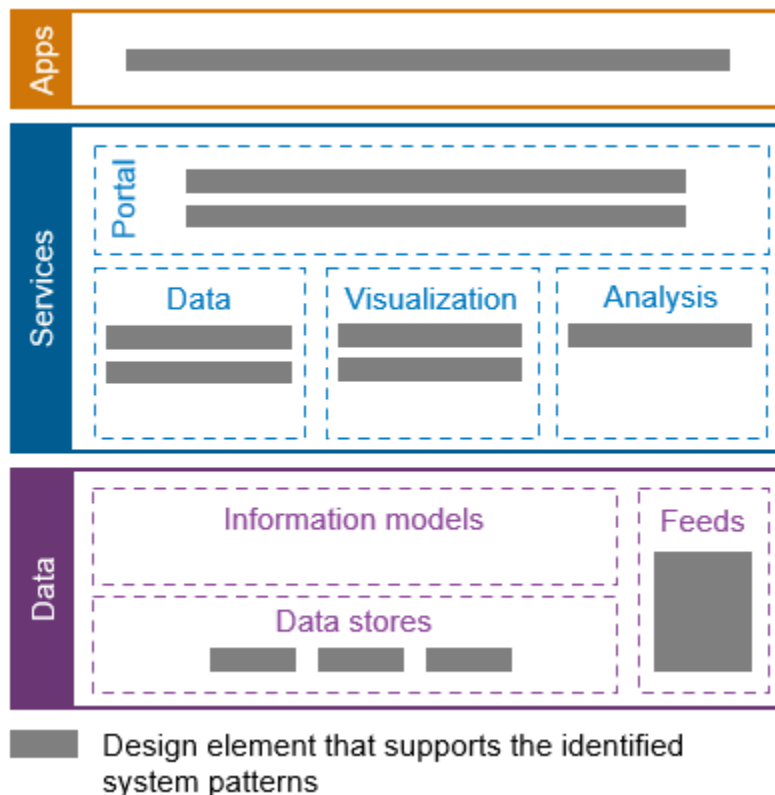


Figure 1.2. The architecture of an ArcGIS system can be represented as a diagram that consists of three tiers. When designing the architecture of an ArcGIS system, design elements that support the identified system patterns are added to the diagram.

Architecting an ArcGIS system (continued)

Architectural pillars for a well-architected system

Implementing the appropriate system patterns to enable the capabilities that your organization needs is necessary but not sufficient for designing a well-architected deployment of ArcGIS Enterprise. You also need to make sure that the system can deliver those capabilities effectively. Six architecture pillars represent IT best practices for implementing a well-architected system:

- Automation
- Integration
- Observability
- Performance and scalability
- Reliability
- Security



[ArcGIS Architecture Center Help](#)

Explore the ArcGIS system patterns

In this activity, you will further explore the eight ArcGIS system patterns that support business needs.

Instructions

- a** Navigate to the [ArcGIS Architecture Center](https://architecture.arcgis.com/en/) (https://architecture.arcgis.com/en/).
- b** At the top of the page, click Patterns.
- c** Review the system patterns and choose a pattern that interests you.
- d** Answer the first question in your workbook.
- e** Click the link for the system pattern of your choosing.
- f** On the right side of the page, click Architecture Considerations.
- g** Review the information in the Architecture Considerations section and answer the second question.
- h** After you are finished, close the Architecture Center browser tab.

1. What business need does the ArcGIS system pattern address?

| ArcGIS system pattern | Business need |
|-----------------------|---------------|
| | |
| | |
| | |
| | |
| | |

Table continued on next page

Explore the ArcGIS system patterns (continued)

| ArcGIS system pattern | Business need |
|-----------------------|---------------|
| | |
| | |
| | |

2. What are two important architectural considerations for this system pattern?

Providing capabilities to support system patterns

Recall that ArcGIS system patterns are common systems that can be built with ArcGIS. Many of these systems rely on additional ArcGIS Server components that are licensed with server licensing roles to provide the required capabilities.

The ArcGIS GIS Server role is the foundational server role for ArcGIS Enterprise as it brings GIS resources, such as map services, feature services, and geoprocessing services, to users as layers in your organization. It is also the licensing role that is needed to deploy a hosting server in your base ArcGIS Enterprise deployment. Organizations may consider deploying additional GIS Servers to handle a greater volume of requests to their services.

In addition to the GIS Server role, there are eight server licensing roles that provide additional capabilities. Your business needs determine which licensing roles that you will use. For example, the GeoEvent Server role supports a comprehensive set of GIS workflows applied to streaming real-time data.

There are several server licensing roles that are available:

- ArcGIS GIS Server
- ArcGIS Image Server
- ArcGIS GeoEvent Server
- ArcGIS GeoAnalytics Server
- ArcGIS Notebook Server
- ArcGIS Knowledge Server
- ArcGIS Workflow Manager Server
- ArcGIS Video Server

Providing capabilities to support system patterns (continued)

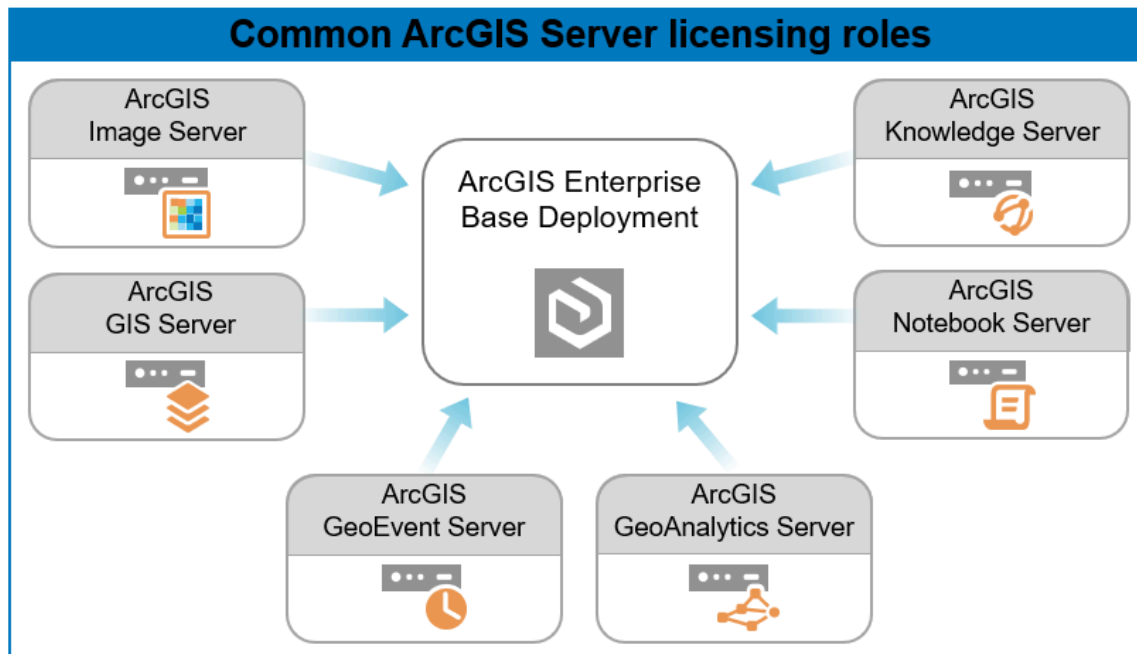


Figure 1.3. ArcGIS Server licensing roles each provide different capabilities and have different system requirements.

You can have many GIS Servers installed in your deployment that can take on any combination of server licensing roles. Each ArcGIS Server site is generally installed on its own hardware as modularization is a key aspect of the ArcGIS Enterprise design.



ArcGIS Enterprise Help: ArcGIS Server 11.3 system requirements

ArcGIS Enterprise Help: ArcGIS Server licensing roles

Using the ArcGIS Enterprise Functionality Matrix

The ArcGIS Enterprise Functionality Matrix represents a quick reference guide to mapping your business needs to specific software functionality.

| Server Capabilities | ArcGIS GeoAnalytics Server | ArcGIS GeoEvent Server | ArcGIS Image Server | ArcGIS Knowledge Server | ArcGIS Notebook Server [†] | ArcGIS Workflow Manager Server | ArcGIS Video Server |
|---|----------------------------|------------------------|---------------------|-------------------------|-------------------------------------|--------------------------------|---------------------|
| Run on Windows | ● | ● | ● | ● | ● | ● | ● |
| Run on Linux | ● | ● | ● | ● | ● | ● | ● |
| Deploy in the cloud | ● | ● | ● | ● | ● | ● | ● |
| Deploy on-premises | ● | ● | ● | ● | ● | ● | ● |
| Deploy disconnected from the open internet | ● | ● | ● | ● | ● | ● | ● |
| Script and automate workflows | ● | ● | ● | ● | ● | ● | ● |
| Create analytic models and model chains | ● | ● | ● | | ● | | |
| Edit data on the web | ● | ● | ● | ● | ● | ● | |
| Create OGC-compliant web services | | | ● | | | | |
| Convert location information to x,y coordinates (geocode) | ● | | | | ● | | |
| Create dynamic image and raster mosaics | | | ● | | | | |

Figure 1.4. The ArcGIS Enterprise Functionality Matrix identifies the capabilities supported by the various server licensing roles so that you can determine which one best meets your needs.

Using the ArcGIS Enterprise Functionality Matrix (continued)

The following table presents an example set of business needs, the server roles that enable the needed capabilities, and the apps that users will use to interact with ArcGIS Enterprise to complete their work. The ArcGIS Conceptual Reference Architecture in combination with the ArcGIS Enterprise Functionality Matrix can be used to inform strategic planning for your ArcGIS Enterprise deployment.

| | |
|-------------------------------------|---|
| Business needs | <ul style="list-style-type: none">• Analysts share a mosaic dataset as an image service• Analysts share a geoprocessing service as a web tool• Crews create and review profile information• Crews access sync-enabled data to capture data (disconnected) in the field |
| Server roles to enable capabilities | <ul style="list-style-type: none">• Image Server• GIS Server |
| User apps | <ul style="list-style-type: none">• ArcGIS Pro• ArcGIS Web AppBuilder• ArcGIS Field Maps |

Business needs and ArcGIS Server licensing roles

Read each of the following scenarios as a set of business needs for different hypothetical companies or situations. Review the business needs, and then answer the questions.

The following resources may be helpful when answering the questions:

- ArcGIS Enterprise Functionality Matrix (C:\EsriTraining\EADM\OnlineFiles)
- ArcGIS Enterprise Help: *ArcGIS Server licensing roles*

Scenario 1: Regional council of governments

A regional council of governments supplies a large set of orthophotos to its participating agencies. The orthophotos must be stored as mosaic datasets. The council intends to offer the mosaics as an imagery layer along with raster analytical tools as the agencies do not have the infrastructure to download and process the imagery themselves.

1. What business needs can you identify?

2. Which additional ArcGIS Server licensing role would meet the organization's needs?

Business needs and ArcGIS Server licensing roles (continued)

Scenario 2: Electrical repair company

A national electrical repair company has many field technicians. The company would like to track its vehicles in real time. Additionally, the company wants to trigger alerts to technicians when an incoming service request is within a certain distance of their location. The company also wants the ability to filter by location, as well as by other attributes (like technician skill sets).

Business needs:

- Create stream service from real-time data stream.
- Filter data by spatial location or attributes.
- Create geofences to trigger alerts.

3. Which additional ArcGIS Server licensing role would meet the organization's needs?

Business needs and ArcGIS Server licensing roles (continued)

Scenario 3: Community advocacy group

A community advocacy group is assessing bus equity. The group would like to uncover, understand, and visualize the relationships of existing bus stop locations and the communities that they serve through graph analysis. They need to incorporate spatial data such as bus stops, routes, and important community locations such as grocery stores and schools. However, they also will need to leverage nonspatial data such as demographic data associated with neighborhoods and bus schedules.

Business needs:

- Share and visualize relationships with a graph service.
- Leverage both spatial and nonspatial data.
- Store analysis results in graph database.

4. Which additional ArcGIS Server licensing role would meet the organization's needs?

Lesson review



ArcGIS Enterprise Help: ArcGIS common patterns of use

1. Before implementing your ArcGIS Enterprise deployment, it is crucial to thoroughly assess your business needs.
 - a. True
 - b. False

 2. What three tiers make up the architecture of an ArcGIS system?
-

3. Which server licensing role provides capabilities that support the location services ArcGIS system pattern?
-

Answers to lesson 1 questions

Explore the ArcGIS system patterns (page 1-5)

1. What business need does the ArcGIS system pattern address?
2. What are two important architectural considerations for this system pattern?

Answers will vary based on the selected ArcGIS system pattern.

Business needs and ArcGIS Server licensing roles (page 1-11)

Scenario 1: Regional council of governments

1. What business needs can you identify?
 - **Storage for a large set of orthophotos as a mosaic dataset**
 - **Support image services**
 - **Raster analytics capabilities**
2. Which additional ArcGIS Server licensing role would meet the organization's needs?

ArcGIS Image Server

Scenario 2: Electrical repair company

3. Which additional ArcGIS Server licensing role would meet the organization's needs?

ArcGIS GeoEvent Server

Scenario 3: Community advocacy group

4. Which additional ArcGIS Server licensing role would meet the organization's needs?

ArcGIS Knowledge Server

2

Adding capabilities through federation

You may initially implement ArcGIS Enterprise as a base deployment. However, over time, you are likely to require more capabilities, or you may simply need to scale out from the initial configuration. Doing so almost certainly involves adding additional ArcGIS Server machines, no matter which server licensing roles that you choose. You may also want to implement a highly available ArcGIS Enterprise deployment, which would require that additional server sites be configured.

As you add additional ArcGIS Server sites, it is important to integrate them with the Enterprise portal to maintain a unified security and sharing model. This integration will decrease administrator burden by having just one security model and one set of workflows for sharing and access by Enterprise portal members.

Topics covered

- Review of the base ArcGIS Enterprise deployment

- Integrating additional Enterprise capabilities into your deployment

- Federating an additional ArcGIS Server site

Base ArcGIS Enterprise deployment

At the core of ArcGIS Enterprise is the base ArcGIS Enterprise deployment, which provides the essential framework that supports a comprehensive GIS.

Using the following diagram, determine the appropriate components and their basic functions within the ArcGIS Enterprise base deployment.




Figure 2.1. The base ArcGIS Enterprise deployment components are shown in this diagram.

1. Component A

2. Component B

Base ArcGIS Enterprise deployment (continued)

3. Component C

4. Component D

5. How does the process of federation affect an ArcGIS Server site?

Federating additional server sites

Over time, you will want to build upon the base deployment configuration for many reasons, such as to expand capabilities, optimize performance, or provide redundancy of the system. Federation is a key aspect of scaling your base ArcGIS Enterprise deployment. Federation applies a systems composition approach for realizing a single ArcGIS system that addresses the ArcGIS system patterns that are required by your organization.

Federation is the method by which ArcGIS Server sites are fully integrated into an ArcGIS Enterprise deployment. By federating a server site, the identity-based security model that the Enterprise portal uses replaces the built-in role-based security model of the server site. There are various benefits to federation, but a core capability is the ability to host tile layers, imagery layers, feature layers, and scene layers that are published by members of the portal.

ArcGIS Image Server and federation

Each server licensing role has unique considerations when adding additional ArcGIS Server sites to your base ArcGIS Enterprise deployment. You will explore several server licensing roles throughout the course, but you will now explore ArcGIS Image Server as an example of adding an additional server licensing role.

ArcGIS Image Server contains several capabilities:

| | |
|------------------------|---|
| Dynamic image services | <ul style="list-style-type: none">• Provides the ability to make large image collections available through dynamic image services• Dynamic image services are served from individual images or mosaic datasets• Dynamic image services have processing applied on the fly as the data is accessed using desktop, web, and mobile applications |
|------------------------|---|

Table continued on next page

Federating additional server sites (continued)

| | |
|-------------------------|---|
| Imagery hosting server | <ul style="list-style-type: none"> • Provides ability to create image services from mosaic datasets • Provides ability to upload and publish an image collection (ZIP file) as a hosted imagery layer • Requires installation of ArcGIS Server on a target machine and then licensed with Image Server role • Image Server machine must be federated with the ArcGIS Enterprise portal |
| Raster analytics server | <ul style="list-style-type: none"> • Provides scalable distributed processing for large images and raster collections • When configured, a set of raster analysis tools is available within ArcGIS Pro and ArcGIS Enterprise portal Map Viewer • Requires that a raster data store be registered in ArcGIS Server Manager • ArcGIS Server site must be licensed with Image Server and federated to ArcGIS Enterprise portal |
| Ortho mapping | <ul style="list-style-type: none"> • Provides server-side processing of imagery into digital terrain models and orthomosaics • The ortho mapping tools are optimized to enable distributed processing and storage |

The following graphic illustrates a stand-alone Image Server site that has not been federated with ArcGIS Enterprise.

Federating additional server sites (continued)

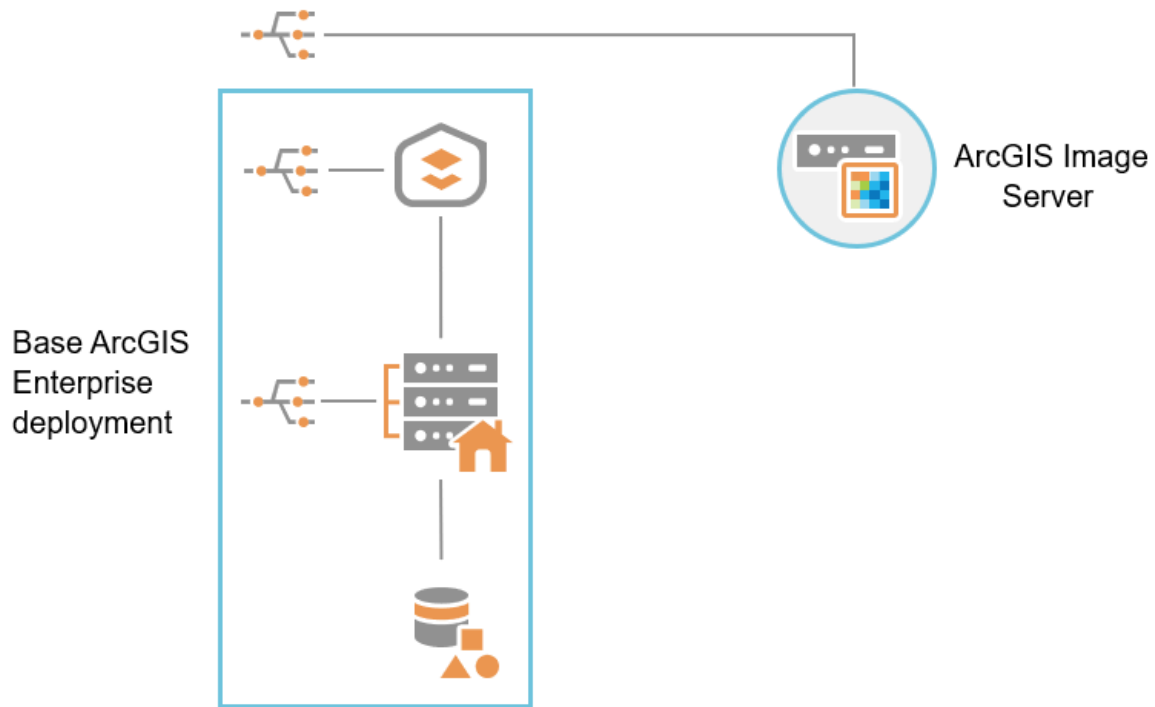


Figure 2.2. An Image Server site can be configured in a stand-alone fashion, but federating it with ArcGIS Enterprise will enable all of its capabilities.

The following graphic shows a minimum configuration of Image Server with both the image hosting and raster analytics capabilities configured on one dedicated Image Server machine.

Federating additional server sites (continued)

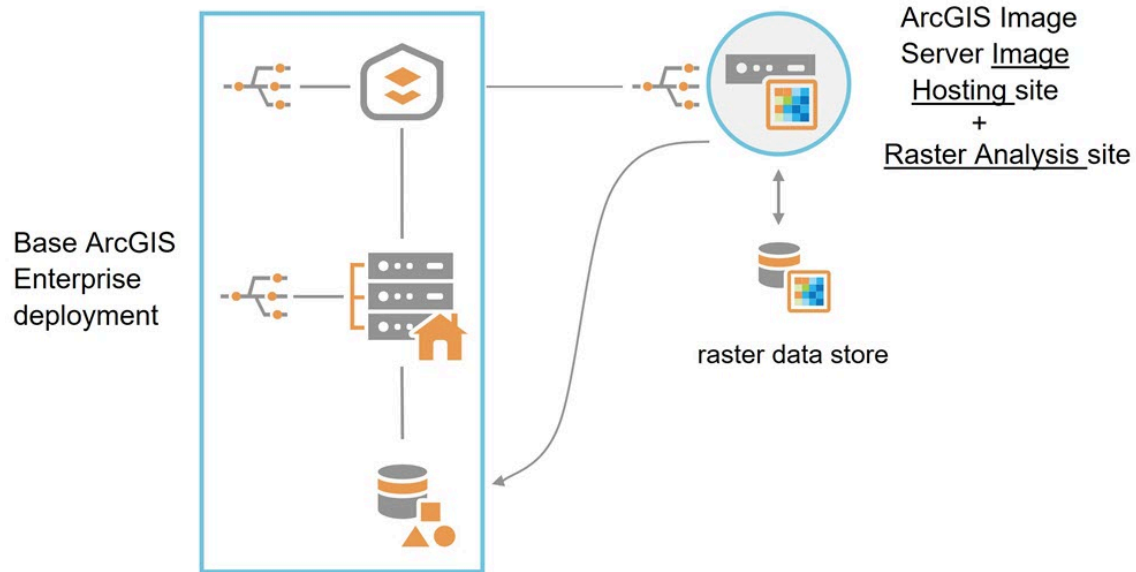


Figure 2.3. This example shows the most basic configuration for a federated Image Server site with both image hosting and raster analytics.



ArcGIS Enterprise Help: *Key concepts for image services*

ArcGIS Enterprise Help: *Configure and deploy raster analytics*

ArcGIS Architecture Center Help: *Using system patterns > Working with multiple system patterns (system composition approach)*

Explore considerations for federating additional server sites

The basic steps for federating a server site are the same whether that server site is to become your hosting server or is an additional server site using any server licensing role.

You will use ArcGIS Enterprise Help documentation to explore several common considerations for federating server sites.

Instructions

- a** Open a web browser and go to the [Portal For ArcGIS: Federate An ArcGIS Server Site With Your Portal](https://enterprise.arcgis.com/en/portal/11.3/administer/windows/federate-an-arcgis-server-site-with-your-portal.htm) web page (https://enterprise.arcgis.com/en/portal/11.3/administer/windows/federate-an-arcgis-server-site-with-your-portal.htm).
- b** Review this page and linked topics to answer the following questions in your workbook.
- c** After you are finished, close the web browser.

1. What are the version requirements for the elements of your base ArcGIS Enterprise deployment?

2. How are existing services on an ArcGIS Server site affected when the site is federated?

Explore considerations for federating additional server sites (continued)

3. After federating, why might an administrator want to disable the ArcGIS Server primary site administrator account (PSA) after the site is federated?

Exercise environment

Throughout this course, you will use an environment that simulates an example ArcGIS Enterprise deployment that consists of multiple virtual machines running in an Azure host PC. The environment will use many of the server licensing roles available for ArcGIS Server. To access the capabilities enabled by these licensing roles, the ArcGIS Server sites must be federated.

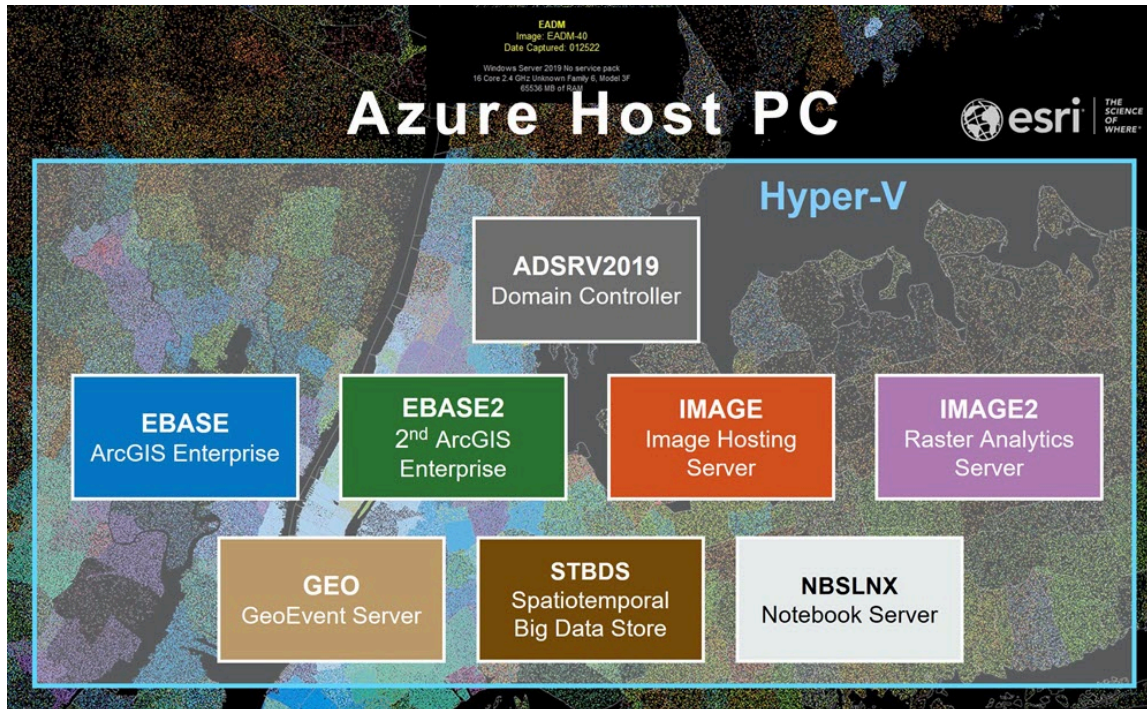


Figure 2.4. This graphic explains the exercise environment that will be used throughout the course.

| | |
|------------------|---|
| ADSRV2019 | Windows Server 2019 domain controller |
| EBASE | Primary ArcGIS Enterprise portal (a base ArcGIS Enterprise deployment) |
| EBASE2 | Secondary ArcGIS Enterprise portal that will be used in the distributed collaboration exercises |
| IMAGE | ArcGIS Image Server machine that will be federated with the Enterprise portal |
| IMAGE2 | Another ArcGIS Image Server machine that will act as the Raster Analytics server |

Table continued on next page

Exercise environment (continued)

| | |
|---------------|---|
| GEO | ArcGIS GeoEvent Server site |
| STBDS | A dedicated spatiotemporal big data store host |
| NBSLNX | A Notebook Server installation on a Linux machine |

Throughout this course, you will use the capabilities provided by multiple licensing roles. To access these capabilities, you will federate the IMAGE, IMAGE2, GEO, and NBSLNX server machines with the base ArcGIS Enterprise deployment on the EBASE machine during the exercises.

Exercise 2

Federate an Image Server site

It is important for administrators to know how to federate an ArcGIS Server site with an ArcGIS Enterprise portal. In this exercise, imagine that your organization currently has a base ArcGIS Enterprise deployment. You have decided to extend the base deployment by installing ArcGIS Server on a dedicated server and assigning the site to the ArcGIS Image Server licensing role. The Image Server site will then be configured for dynamic image services and raster analytics.

You will first federate an Image Server site with the ArcGIS Enterprise portal and then verify that the operation has been successful. Additionally, you will use the raster analysis capabilities of the federated Image Server site, which is available to ArcGIS Pro and Map Viewer users.

The following operations have already been performed on the Image Server site:


- Installed ArcGIS Server
- Licensed with the ArcGIS Image Server licensing role
- Installed ArcGIS Web Adaptor with the name "image" for the ArcGIS Image Server
- Configured SSL with a domain certificate


In this exercise, you will perform the following tasks:

- Federate an Image Server site with the ArcGIS Enterprise portal.
- Verify the status of federation.
- Explore the raster analysis tools.

Step 1: Examine a stand-alone Image Server site

You will begin by investigating a stand-alone Image Server site, which is not yet federated with the Enterprise portal. You will first connect to the IMAGE machine to explore the current state of the ArcGIS Image Server configuration.

- a On the host VM, from the Windows Taskbar, click the Remote Desktop Connection Manager icon .
- b On the left side, in the server tree panel, under EADM_Sessions, double-click IMAGE to connect to the IMAGE VM.



- c Sign in using the adAdmin account (password: **Esri.4.GIS**), if necessary.



The desktop to the IMAGE VM is orange.

- d Using the connection to the IMAGE VM, open a web browser.

You will use ArcGIS Server Manager to explore your Image Server site before federation.

- e On the web browser, click the IMAGE-Server Manager bookmark (<https://image.ad.local/image/manager>), if necessary.

This site is the ArcGIS Server Manager application running on the IMAGE machine.

- f Sign in using the PSA (Primary Site Administrator) credentials (username: **siteadmin**; password: **Esri.4.GIS**).



For convenience when signing in throughout the exercises in this course, credentials have been saved in the web browser. You may use the saved credentials or type the username and password.

- g In ArcGIS Server Manager, click the Security tab.

- h Review the Configuration Settings section.

1. What User Store and Role Store are used by the site?

- i Click the Services tab.
- j Notice the ElevationAthensTX image service.


This image service has been published from a National Elevation Dataset (NED) raster.

- k Click the ElevationAthensTX image service.
- l On the left side, click the Capabilities tab.
- m Under URLs, notice the URLs that are used to access the service.

When you federate the Image Server site, the web layer that is created in the Enterprise portal will reference the REST URL.

Step 2: Investigate a raster store

In addition to the image hosting capabilities of Image Server, you want to use the raster analytics capability to enable distributed processing and storage for raster analysis. Registering a raster store is a prerequisite to using these capabilities. This raster store is used for storing hosted raster datasets that are generated by raster analysis tools and the rasters that portal users upload when creating imagery layers using the ArcGIS Enterprise portal.

- a In ArcGIS Server Manager, click the Site tab.
- b On the left side, click the Data Stores tab.
- c In the table, under Type, for the Raster Store data store type, click the Edit button .

2. What is the file path for this data store?

A shared folder on the IMAGE machine has been registered as a raster store so that other machines in ArcGIS Enterprise will be able to access the data written there.



This raster store is specified using a UNC path, rather than a local file path (**C:\RasterAnalyticsOutput**).

- d Click Cancel to close the Edit Raster Store dialog box.

Additional raster stores may be beneficial for implementing redundancy, increased performance, or workload separation. Those raster stores can be configured on this page in ArcGIS Server Manager.

Step 3: Federate the Image Server site

Now you will federate the Image Server site with your ArcGIS Enterprise deployment.

- a Using the Remote Desktop Connection Manager, in the server tree panel, double-click EBASE to connect to the EBASE VM.
- b Sign in using the adAdmin account (password: **Esri.4.GIS**), if necessary.



The desktop to the EBASE VM is blue.

- c In File Explorer, browse to **\\EADM\EsriTraining\EADM\FedEnablesCaps** and notice the FederateImageServer.txt file.

This file contains the URL that you will copy to ensure accuracy as you federate the server.

- d Open the FederateImageServer.txt file and leave it open for the remainder of the step.
- e Open a web browser and click the bookmark for ArcGIS Enterprise, if necessary.

You will now sign in as an administrative user.

- f Click Sign In.
- g Sign in to the Enterprise portal with the **portaladmin** username and **Esri.4.GIS** password.
- h In the Enterprise portal, near the top of the page, click Content.
- i On the Content page, click the My Organization tab.

You want to search the entire organization, rather than only your own content. Because you are signed in as an administrator, you have access to items that are owned by any member of the

portal.

j In the Search My Organization field, type **ElevationAthensTX** and press Enter.

The search will not return any results because the ArcGIS Server site has not yet been federated with the ArcGIS Enterprise portal.

k Near the top of the page, click Organization, and then click the Settings tab to access the ArcGIS Enterprise settings.

l On the Settings page, on the left side, click the Servers tab.

m Under Federated Server Sites, click Add Server Site, and then specify the following information:

- Services URL: **https://image.ad.local/image** (copy and paste from the text file)
- Administration URL: **https://image.ad.local:6443/arcgis** (copy and paste from the text file)
- Username: **siteadmin**
- Password: **Esri.4.GIS**


n Click Next.



Federating the server may take several minutes.

When the server is federated, you will advance to the Configure Server Role page. This page is where you will designate this Image Server site as the Raster Analysis Server for your organization.

o Next to Raster Analysis Server, enable the option to configure this server site with this role.



p Click Save Server Role.



It may take a moment for the server to federate.

- q At the bottom of the list of federated server sites, notice the federated server that appears.

| | |
|---------------------|---|
| image.ad.local:6443 | ... |
| Status | ✓ All systems operational |
| Server role | Raster Analysis Server |
| Service URL | https://image.ad.local/image |
| Administration URL | https://image.ad.local:6443/arcgis |



In certain circumstances, it can be advantageous to use the web adaptor URL to specify both the service URL and the administration URL. For example, if you are federating a multi-machine or highly available ArcGIS Server site, or if your site is hosted in a cloud environment, use the web adaptor URL in the Administration URL field.

- r Next to Server Role, click Raster Analysis Server.

You are returned to the Configure Server Role page, where you could change the server role for this Image Server site.

- s Notice that the Image Hosting Server role is enabled in addition to the Raster Analysis Server role.

Even though you did not choose the Image Hosting Server role when you federated the Image Server site, it was automatically enabled because you have a single Image Server site. When you designate a raster analysis server, it implicitly acts as the image hosting server. With this configuration, this single site must perform both roles.

- t Click Done.

Step 4: Examine the effects of federation

To verify that the new Image Server has been federated, you will look at the Web Administrator Directory for the Image Server site.

- a Using the Remote Desktop Connection Manager, restore the IMAGE VM (orange desktop).

- b In a web browser, go to **https://image.ad.local/image/admin**.
- c Sign in using the PSA (Primary Site Administrator) credentials (username: **siteadmin**; password: **Esri.4.GIS**).
- d Click the Security link, and then click the Config link.

3. How can you tell that this server site has been federated?

| | |
|-------------------------|-------------------------|
| Server role: | FEDERATED_SERVER |
| Server function: | RasterAnalytics |

-User Store Configuration

| | |
|--------------|---------------|
| Type: | PORTAL |
|--------------|---------------|

-Role Store Configuration

| | |
|--------------|---------------|
| Type: | PORTAL |
|--------------|---------------|

- e Open a new web browser tab and click the ArcGIS Enterprise bookmark.
- f Sign in using the **portaladmin** username and **Esri.4.GIS** password.
- g Click Content.

4. Is there a portal item for the ElevationAthensTX service?



- h Click ElevationAthensTX.
- i On the item details page, scroll down and locate the Owner information for this item.


5. Who owns this item?



The administrator who performs the federation will be the owner of the portal items that are

created during federation. Reassigning ownership of these items is a common workflow that follows federation.

- j Next to Owner, click Change Owner.
- k In the Change Owner dialog box, click Publisher Built-in.



- l Click Save to change the item's ownership.
- m Verify that Publisher now owns this item.

In this step, you confirmed that the Image Server site was federated and saw how existing services were incorporated into the Enterprise portal.


- n Close the web browser.

Step 5: Use the raster analysis tools


Now you can explore the raster analysis tools that have been enabled for both Map Viewer and ArcGIS Pro.

- a Restore the EBASE VM (blue desktop).
- b Sign out of the ArcGIS Enterprise portal, and then sign in with the **publisher** username and **Esri.4.GIS** password.


You will begin by creating a new content folder to store the results of the slope analysis that you are preparing to do.

- c Click Content, and then next to Folders, click the Create Folder button  .
- d For Folder Name, type **AnalysisResults** and click Save.



You will create a map using the ElevationAthensTX layer, which will be the input for your slope analysis.

- e Near the top of the page, click Map to open a new web map.
- f On the Layers pane, click Add.
- g For the ElevationAthensTX layer, click Add.
- h To the left of Add Layer, click the Back button  to return to the Layers pane.

You will now zoom the map to the area where you want to evaluate slope.

- i In the bottom-right corner of the map, click the Search button .
- j In the Search field, type **Athens, TX** and press Enter.
- k In the Search Result pop-up, click Zoom To one time.
- l Close the Search Result pop-up in the map.

You are now ready to conduct a slope analysis of this elevation (raster) layer.








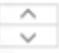




- m On the right side of the map, from the Settings toolbar, click the Analysis button .
- n On the Analysis pane, click Tools.
- o Click the Filter button  and choose Raster Analysis.
- p Expand the Analyze Terrain section and click Surface Parameters to open the tool.

Now that the Surface Parameters tool is open, you will specify the parameters to perform a slope analysis for the current map extent.

- q For Input Surface Raster, click Layer and choose ElevationAthensTX.
- r Under Surface Parameter Settings, verify the following parameters:
 - Parameter Type: Slope
 - Slope Measurement: Degree






Surface parameter settings

Specify the type of surface parameter and related parameters. Define the neighborhood size and whether it will adapt to the terrain.

| | |
|--|---|
| Parameter type |  |
| <input type="text" value="Slope"/> |  |
| Slope measurement |  |
| <input type="text" value="Degree"/> |  |
| Local surface type |  |
| <input type="text" value="Quadratic"/> |  |
| Neighborhood distance |  |
| <input type="text" value="Enter a distance v..."/> |  |
| <input type="text" value="Unit"/> |  |
| <input type="checkbox"/> Use adaptive neighborhood |  |
| Z unit |  |
| <input type="text" value="Select an option..."/> |  |

S Under Result Layer, complete the following parameters:

- Output Raster Name: **SlopeAthens**
- Save In Folder: Analysis Results

| | |
|--|---|
| Result layer | |
| Provide a unique name for the result layer. | |
| Output raster name  |  |
| <input type="text" value="SlopeAthens"/> | |
| Save in folder |  |
|  AnalysisResults |  |

t Expand the Environment Settings section.

u For Processing Extent, choose Display Extent.

v Click Run.



It may take a minute for the analysis to complete.

w Briefly explore the result of your slope analysis.

x Close the web browser without saving the web map.

y Close the FederateImageServer.txt file and close File Explorer.

One final consideration is which Enterprise portal members will be able to use the raster analysis capabilities. By default, the publisher and administrator roles in the Enterprise portal are granted the ability to use raster analysis. However, if you create any custom roles that you want to enable these capabilities for, you will need to explicitly grant permissions in the Enterprise portal.

In this exercise, you federated an Image Server site, configured it as both the image hosting and raster analytics server for your deployment, and tested the raster analysis capabilities that are available in Map Viewer.

Lesson review

1. A raster analysis is run using the raster analytics server. Which component is needed to store the output?
 - a. A data store
 - b. An enterprise geodatabase
 - c. A file geodatabase
 - d. A raster data store

 2. What capabilities are gained by federating an Image Server site with the ArcGIS Enterprise portal?
-
-

Answers to lesson 2 questions

Base ArcGIS Enterprise deployment (page 2-2)

1. Component A

ArcGIS Web Adaptors; web adaptors forward requests to your ArcGIS Server machines. If web-tier authentication is not required, other third-party options that support load balancing or reverse proxy can replace the ArcGIS Web Adaptor component.

2. Component B

Portal for ArcGIS; the ArcGIS Enterprise portal allows you to create, share, and manage maps, scenes, apps, and other geographic information with other people in your organization.

3. Component C

ArcGIS Server; ArcGIS Server is server software that makes your GIS resources available through web services.

4. Component D

ArcGIS Data Store; this component provides data storage for the hosting server that is used with your ArcGIS Enterprise portal.

5. How does the process of federation affect an ArcGIS Server site?

When you federate a server with your portal, the portal's security store controls all access to the server. Services that exist on the ArcGIS Server site at the time of federation are automatically added to the portal as items.

Explore considerations for federating additional server sites (page 2-8)

1. What are the version requirements for the elements of your base ArcGIS Enterprise deployment?

The elements (including the hosting server) must be at the same version as your portal.

Answers to lesson 2 questions (continued)

2. How are existing services on an ArcGIS Server site affected when the site is federated?

Existing services remain, and corresponding items are added to the portal. The items will be owned by the portal administrator who performs the federation, but the ownership of these items can be reassigned as needed after federation.

3. After federating, why might an administrator want to disable the ArcGIS Server primary site administrator account (PSA) after the site is federated?

Disabling the PSA enhances the security of your system by ensuring that the only way to administer ArcGIS Server is through the group or role that you have specified in your identity store, such as your portal's built-in identity store.

Exercise 2: Federate an Image Server site (page 2-12)

1. What User Store and Role Store are used by the site?

ArcGIS Server Built-in

2. What is the file path for this data store?

\\image\RasterAnalyticsOutput

3. How can you tell that this server site has been federated?

The server role is set to FEDERATED_SERVER, and both the User Store Configuration and Role Store Configuration are set to PORTAL.

4. Is there a portal item for the ElevationAthensTX service?

Yes

5. Who owns this item?

portaladmin

3

Ensuring resilient scaling through distributed computing

Administrators of ArcGIS Enterprise have a wide range of objectives to meet, including maintaining high performance, implementing new functionality, and managing access. Applying architectural best practices allows ArcGIS Enterprise administrators to maximize the performance, reliability, and scalability of their system.

ArcGIS Enterprise is well suited to using distributed computing to help meet these objectives. This lesson will focus on several best practices for using distributed computing to support an evolving ArcGIS Enterprise deployment.

Topics covered

- Distributed computing advantages

- High availability in ArcGIS Enterprise

- Practicing workload separation in an evolving deployment

Distributed computing and business needs

You can use software components installed across multiple computers that work together to achieve business goals. This concept is called distributed computing, and it is foundational to ArcGIS Enterprise. For example, a base ArcGIS Enterprise deployment exemplifies distributed computing when its constituent software components are installed on dedicated machines.




Figure 3.1. In this example, the four components of a base ArcGIS Enterprise deployment are installed on separate machines, yet they work together to accomplish computing tasks.

Distributed computing concepts and best practices can be incorporated into an ArcGIS Enterprise deployment to better meet various business needs. Consider the following examples in the table that demonstrate distributed computing concepts.

| Distributed computing setup | Business need |
|---|--|
| Multi-machine ArcGIS Server site that shares server directories and a configuration store | Increase the processing power of the site |
| Maintaining at least two copies of your deployment and implementing failover mechanisms | Support the uptime requirements of a service level agreement, even if the primary deployment fails |
| Separating the infrastructure for functions such as data storage, computationally intensive tasks, and analysis functions | Increase the efficiency of resource usage while maintaining capabilities for users |
| Enabling distributed analysis | Support large vector and raster processing operations |

Table continued on next page

Distributed computing and business needs (continued)

| Distributed computing setup | Business need |
|--|--|
| Using a multi-machine implementation of the spatiotemporal big data store in ArcGIS Data Store | Store high-volume observation data |
| Load balancing to distribute traffic to back-end servers | Support an increased demand for web services |

What are some examples of how distributed computing can be used in your organization?

Explore ArcGIS Enterprise high availability

A crucial consideration for any organization implementing ArcGIS Enterprise is to determine what level of operational performance over a given time interval will be required to meet business needs. To maximize uptime, you can create a highly available ArcGIS Enterprise deployment by reducing single points of failure in your system.

High availability is a design approach that supports the reliability pillar of the ArcGIS Well-Architected Framework.

Instructions

- a** On the host machine, navigate to the [ArcGIS Architecture Center](https://architecture.arcgis.com/en/) (<https://architecture.arcgis.com/en/>).
- b** At the top of the page, click Pillars.
- c** Under Reliability, click High Availability.
- d** Read the Define And Understand Availability section and the Design Considerations section, and then answer the following questions in your workbook.
- e** After you are finished, close the current browser tab.

1. What is the difference between high availability and disaster recovery?

2. What are two metrics that an organization might use to measure whether the system is meeting the necessary level of operational performance over time?

Explore ArcGIS Enterprise high availability (continued)

3. What are six design constraints that can affect how high availability is achieved?

Making ArcGIS Enterprise components highly available

Incorporating high availability design approaches in an ArcGIS Enterprise deployment improves the reliability of the system. High availability can be applied to all components of ArcGIS Enterprise or may only be applied to certain components in the deployment based on business needs. Configuring high availability is different for each component of ArcGIS Enterprise.

Portal for ArcGIS

An ArcGIS Enterprise portal is made highly available by configuring two Portal for ArcGIS machines that are accessed through a load balancer. One machine functions as the primary node, and the other machine functions as a standby node. If the primary machine becomes unavailable, the standby machine is promoted to primary. When the unavailable portal is restarted or becomes available, it rejoins as the standby node. The primary and standby nodes share state through inter-machine communication via ports and database synchronization but also rely on shared file storage for the portal's content directory.



For your portal to be highly available, you must configure the portal content directory on a highly available file server.

Overall, the highly available portal system is active-passive because only the primary node handles state changes, such as adding new members or content. However, at the web server level, the portal system is active-active because both nodes can service incoming requests and search indices in the shared content directory.

Making ArcGIS Enterprise components highly available (continued)

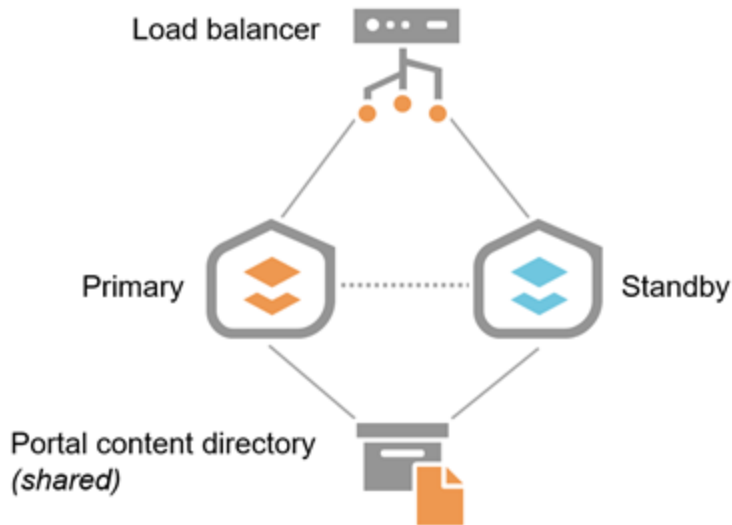


Figure 3.2. In this example, a third-party load balancer forwards requests to the Portal for ArcGIS machines, depending on which one is the primary. Both of the Portal for ArcGIS components share a content directory.

ArcGIS Server

A highly available ArcGIS Server site in an active-active configuration consists of two or more fully redundant machines joined to the site with a load balancer that routes requests to every machine. If one node becomes unavailable, requests are sent to any of the other nodes that participate in the site. The nodes share a configuration store and server directories.



For your ArcGIS Server site to be highly available, you must configure the configuration store and server directories on a highly available file server.

Making ArcGIS Enterprise components highly available (continued)

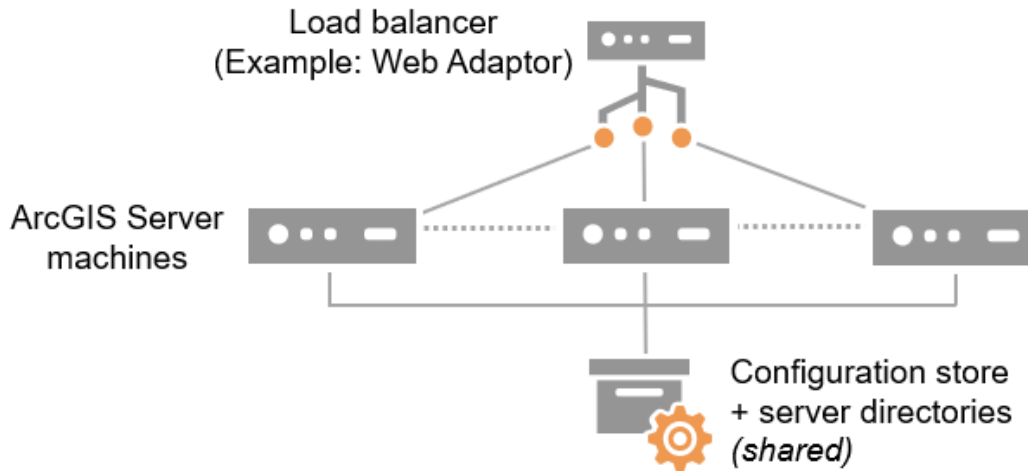


Figure 3.3. In this example, a third-party load balancer forwards requests to three ArcGIS Server machines participating in a multi-machine site that share a configuration store and server directories.

ArcGIS Web Adaptor

If your portal is using web-tier authentication, you will need to configure ArcGIS Web Adaptor instances for each portal in addition to a third-party load balancer, as shown in the following diagram. In this case, the ArcGIS Web Adaptor instances are fully redundant in an active-active configuration.

Making ArcGIS Enterprise components highly available (continued)

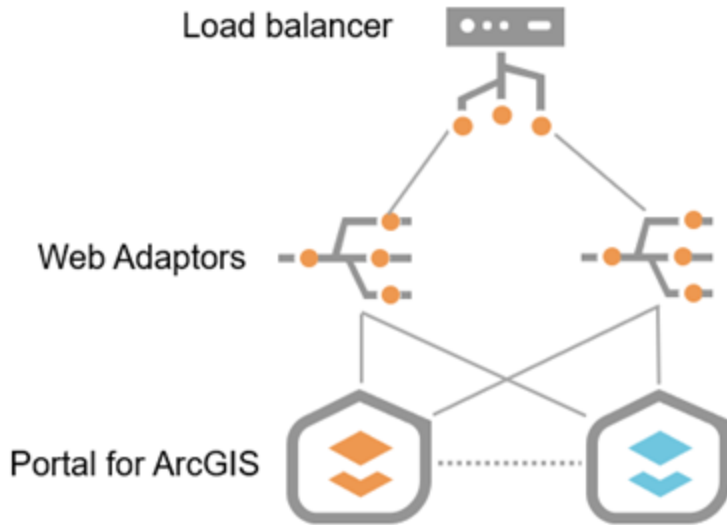


Figure 3.4. In this example, a third-party load balancer forwards requests to Web Adaptors configured for the Portal for ArcGIS components.

Similar to Portal for ArcGIS, you can establish this redundancy at the web server tier for an ArcGIS Server site by installing the Web Adaptor on multiple web servers.

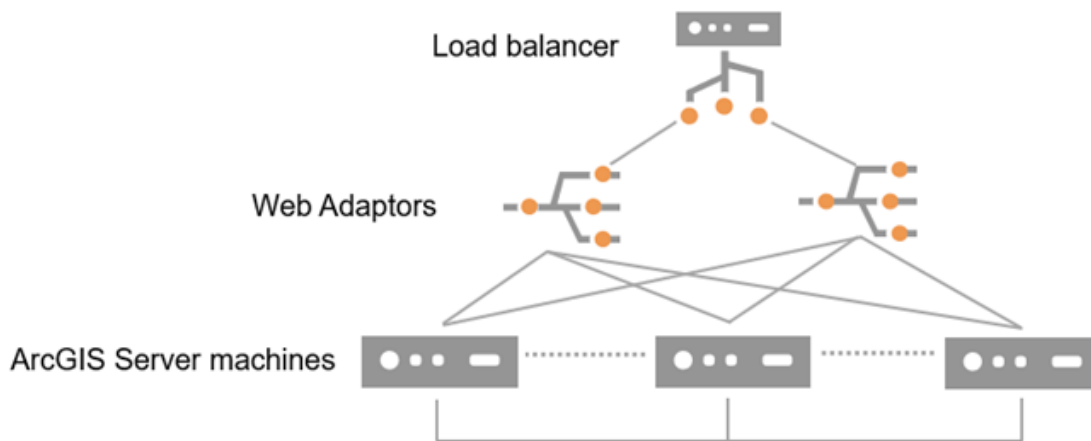


Figure 3.5. In this example, a third-party load balancer forwards requests to two Web Adaptors configured for a multi-machine ArcGIS Server site.

Making ArcGIS Enterprise components highly available (continued)

ArcGIS Data Store

To configure a relational data store to be highly available with an active-passive configuration, a standby instance is installed on a second machine. Feature layer data is replicated from the primary machine to the standby machine. When the primary machine becomes unavailable, the standby machine is promoted to primary.

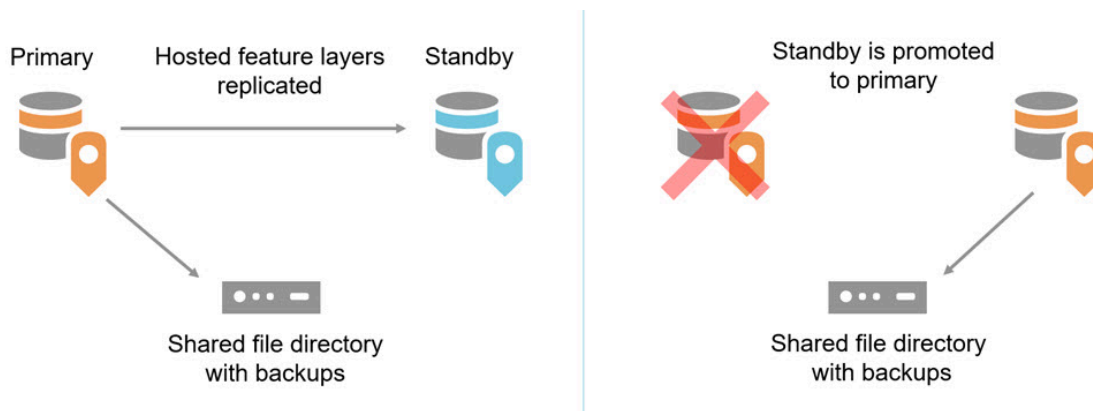


Figure 3.6. This diagram shows the primary machine becoming unavailable and the standby machine being promoted to primary. When the standby becomes the primary relational data store, it sends its backup files to the same location as the former primary machine.

Similar to the relational data store, a tile cache data store can be configured to be highly available with an active-passive system (primary-standby mode). However, tile cache data stores can also be configured with an active-active system (clustered mode), in which scene layers are distributed across multiple machines. In clustered mode, the tile cache data store should contain an odd number of machines with a minimum of three machines.

Making ArcGIS Enterprise components highly available (continued)

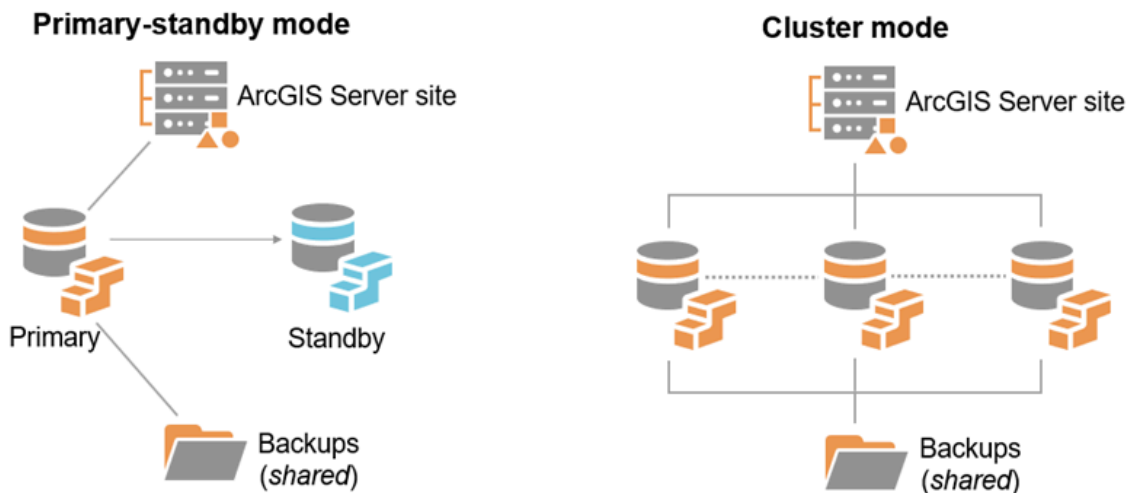


Figure 3.7. These two diagrams illustrate the options for making a tile cache data store highly available.

ArcGIS Enterprise on Kubernetes

High availability is an architectural best practice that can be applied to ArcGIS Enterprise deployed on Linux or Windows. However, there is now a third option to deploy: ArcGIS Enterprise on Kubernetes. An important feature of ArcGIS Enterprise on Kubernetes is that it is innately highly available. Due to its unique architecture, Kubernetes systematically detects when its deployment components are unhealthy or unavailable and will respond by recreating components within the system, maintaining resiliency and availability.



Additional information regarding ArcGIS Enterprise on Kubernetes is outside the scope of this course. If you are interested in a more in-depth discussion of ArcGIS Enterprise on Kubernetes, consider the instructor-led course *Deploying ArcGIS Enterprise on Kubernetes* (KUBE).

Making ArcGIS Enterprise components highly available (continued)

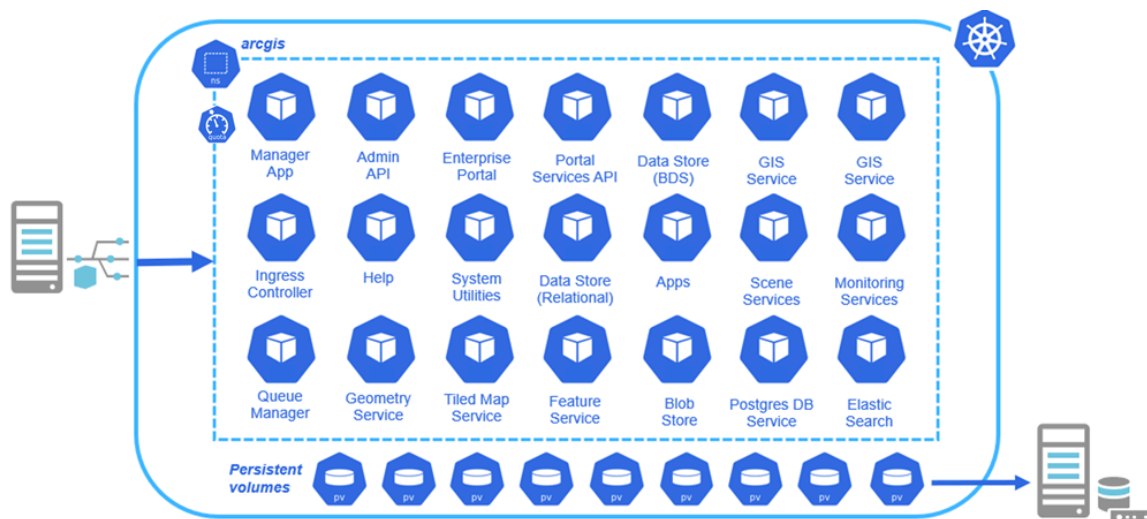


Figure 3.8. This ArcGIS Enterprise on Kubernetes diagram shows ArcGIS Enterprise deployed as a set of microservices based on operational components.

 **ArcGIS Enterprise Help: Configure highly available ArcGIS Enterprise**

Explore workload separation

Workload separation is a design approach that supports the performance and scalability pillar of the ArcGIS Well-Architected Framework. This distributed computing strategy can be used to allocate workloads to appropriate server resources.

Instructions

- a On the host machine, navigate to the [ArcGIS Architecture Center](https://architecture.arcgis.com/en/) (<https://architecture.arcgis.com/en/>).
- b At the top of the page, click Pillars.
- c Under Performance And Scalability, click Overview.
- d On the left, in the Performance And Scalability section, expand Focused Topics.
- e Under Focused Topics, click Workload Separation.
- f Review the documentation, and then answer the following questions in your workbook.
- g After you are finished, close the current browser tab.

1. What are two common approaches to workload separation?

2. Why might you want to isolate visualization and analysis tasks in your deployment?

Comparing high availability and workload separation

Both high availability and workload separation are examples of distributed computing strategies that can be used to meet your organization's business needs. Understanding the differences between the two will help you decide whether it is appropriate to implement one or both strategies.

Review the following two diagrams that provide visual examples of high availability and workload separation. Then, complete the subsequent table below by comparing the characteristics of high availability and workload separation. For each characteristic, identify whether it corresponds to high availability, workload separation, or both.

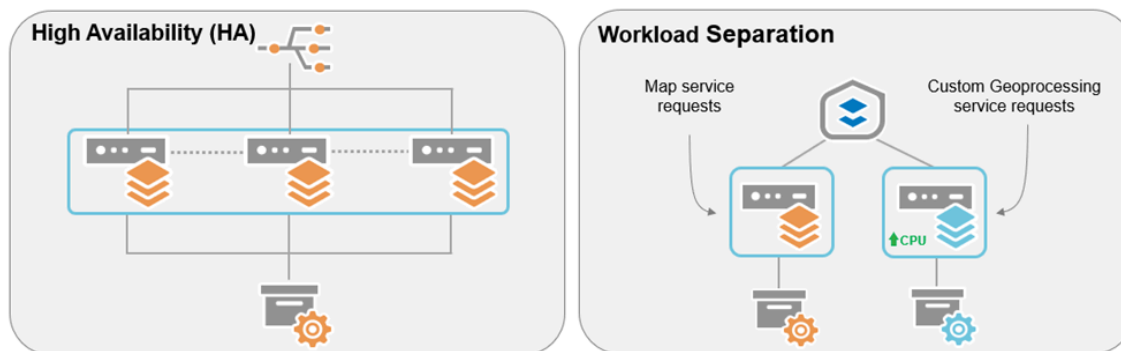


Figure 3.9. The High Availability diagram illustrates a multi-machine ArcGIS Server site, with three participating servers. A load balancer distributes requests to all machines in the site, and they share a configuration store and server directories. The Workload Separation diagram illustrates two ArcGIS Server sites federated with an ArcGIS Enterprise portal. A load balancer distributes requests to the appropriate ArcGIS Server based on the type of request.

| Characteristic | Distributed computing strategy |
|--|--------------------------------|
| Computationally intensive tasks cannot negatively impact other functions | |
| Eliminates a single point of failure through redundant infrastructure | |
| Used to meet business-defined service level agreements (SLAs) | |

Workload separation by server capability

You will continue focusing on the Image Server licensing role as an example of workload separation best practice. However, remember that workload separation is a general concept that applies across all ArcGIS Enterprise server roles. It also highlights the modular nature of the components used to implement a given system infrastructure. It can work together with system scalability to support a robust and distributed GIS.

Image Server

By implementing an Image Server site as a separate machine along with the machines used in your base ArcGIS Enterprise deployment, you are following the best practice of workflow separation.

In addition to hosting tiled and dynamic image services based on scalable mosaic datasets, ArcGIS Image Server supports two specialized capabilities: raster analytics and image hosting. These two server functions represent inherently different types of demands placed on the server; the former is more computationally intensive while the latter is more related to visualizing services. Implementing each server function on its own dedicated ArcGIS Server site is an excellent example of workload separation. This configuration isolates the computational demand of raster analysis from day-to-day imagery support tasks like providing tiled basemaps and dynamic image service products.

Workload separation by server capability (continued)

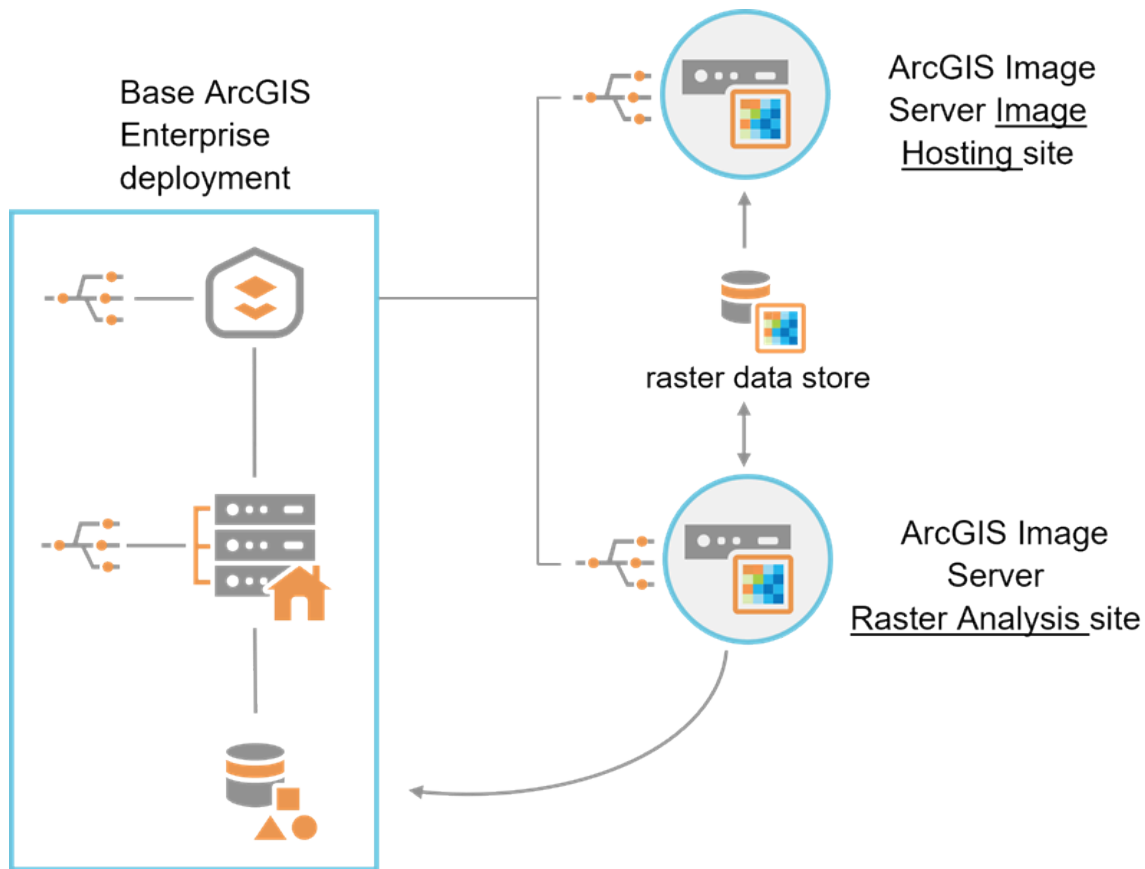


Figure 3.10. The workload separation best practice is illustrated by separating the image hosting and raster analytics capabilities of ArcGIS Image Server.

Exercise 3

Implement workload separation

In an earlier exercise, you configured a federated ArcGIS Image Server site to act as both an image hosting and raster analytics server. Although this configuration worked well initially, users are now reporting unacceptable response times when performing analyses on large sets of rasters. You also discern that the SLA for uptime relating to raster visualization processes has not been met. After further investigation, you determine that the timeouts for raster visualization services coincide with the running of raster analysis tasks.

Therefore, you decide to implement workload separation for Image Server. Your strategy is to dedicate separate hardware resources that are appropriate for performing raster analyses from those performing raster visualization. To test the configuration, you will execute a raster analysis task using ArcGIS Pro and then promptly view ArcGIS Server Manager to review the activity of the associated services.

In this exercise, you will perform the following tasks:


- Federate a second ArcGIS Image Server site.
- Explore raster-oriented system services.
- Use raster analysis tools.

Step 1: Configure a raster store for second Image Server site

First, you will configure your raster store (located on the IMAGE machine) to support the output of raster analysis operations from a second ArcGIS Image Server site.

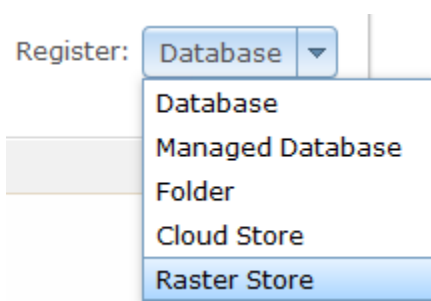
The folder has been shared so that it can be accessed across the network using a UNC path, and the ArcGIS Server account (the account that owns ArcGIS Enterprise services) has been granted permissions to read and write to the folder. Therefore, you are ready to register it in ArcGIS Server Manager as a raster store using the UNC path.

- a Using the Remote Desktop Connection Manager, connect to the IMAGE2 VM.
- b Sign in with the **adAdmin** username and **Esri.4.GIS** password, if necessary.



The IMAGE2 VM has a lavender desktop.

- c Open a web browser and click the bookmark for IMAGE2-Server Manager (https://image2.ad.local/raster_analysis/manager).
- d Sign in with the **siteadmin** username and **Esri.4.GIS** password.
- e Click the Site tab, and then on the left side, click the Data Stores tab.
- f From the Register drop-down list, choose Raster Store, as shown in the following graphic.



You will use File Explorer to verify the correct UNC path to use.

- g Open File Explorer.
- h In the address bar, type the UNC path **\\image\RasterAnalyticsOutput** and press Enter.

The UNC path should resolve, allowing you to view the contents of this folder.

- i Notice that the output from your previous slope analysis is located in this folder.
- j Copy the UNC path from File Explorer.
- k Restore the web browser with the Register Raster Store dialog box open, and then set or verify the following parameters:
 - Name: **RasterStore**
 - Type: File Share
 - Path: **\\image\RasterAnalyticsOutput** (copy and paste from File Explorer to eliminate typos)

Register a raster store on your ArcGIS Server

Name:

RasterStore

Type:

File Share

Path:

\\image\RasterAnalyticsOutput

- l Click Create.
- m Check the box next to RasterStore, and then click Validate.

| Status | File Name | Type | |
|--------|-------------|--------------|---|
| ✓ | RasterStore | Raster Store | ✗ |

- n Click the Services tab, and then on the left side, click the System folder.
- o In the search field, type **raster** and press Enter to expose the raster support services.

1. What is the current state of the raster services?

After federation and setting this machine as the raster analysis server, you will return to see

which services will get started.

- p** Leave the web browser window open for use in a later step.

Step 2: Federate the Image Server site

Now you will federate the second Image Server site with your ArcGIS Enterprise deployment.

- a** Using the Remote Desktop Connection Manager, view the EBASE VM window (blue desktop).

Hint: If you do not have an active connection to the EBASE VM in the Remote Desktop Connection Manager, double-click EBASE to make a connection.

- b** Open File Explorer and browse to `\\EADM\EsriTraining\EADM\UseWorkloadSep`.
- c** Double-click the FederationURLs.txt file.

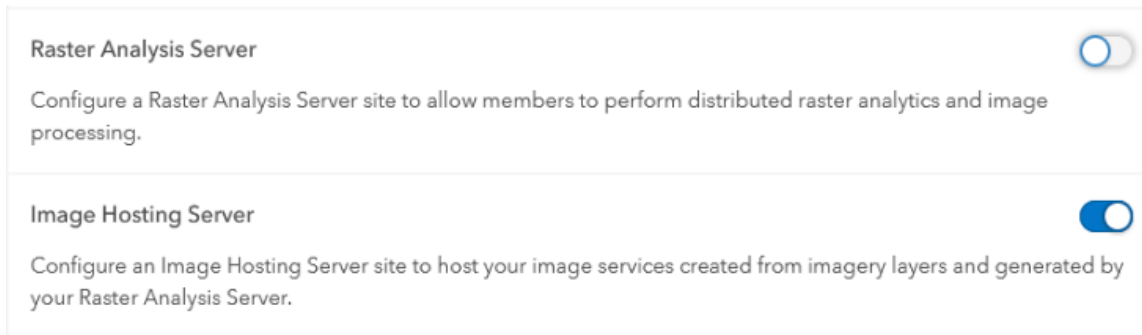
You will use the URLs from this file to eliminate potential typos when federating.

- d** Open a web browser and sign in to the Enterprise portal with the **portaladmin** username and **Esri.4.GIS** password.
- e** Click Organization, if necessary, and then click the Settings tab to access the ArcGIS Enterprise settings page.
- f** On the left side, click the Servers tab.
- g** Click Add Server Site, and then specify the following information:
 - Services URL: **https://image2.ad.local/raster_analysis** (copy and paste from the FederationURLs.txt file)
 - Administration URL: **https://image2.ad.local:6443/arcgis** (copy and paste from the FederationURLs.txt file)
 - Username: **siteadmin**
 - Password: **Esri.4.GIS**
- h** Click Next.

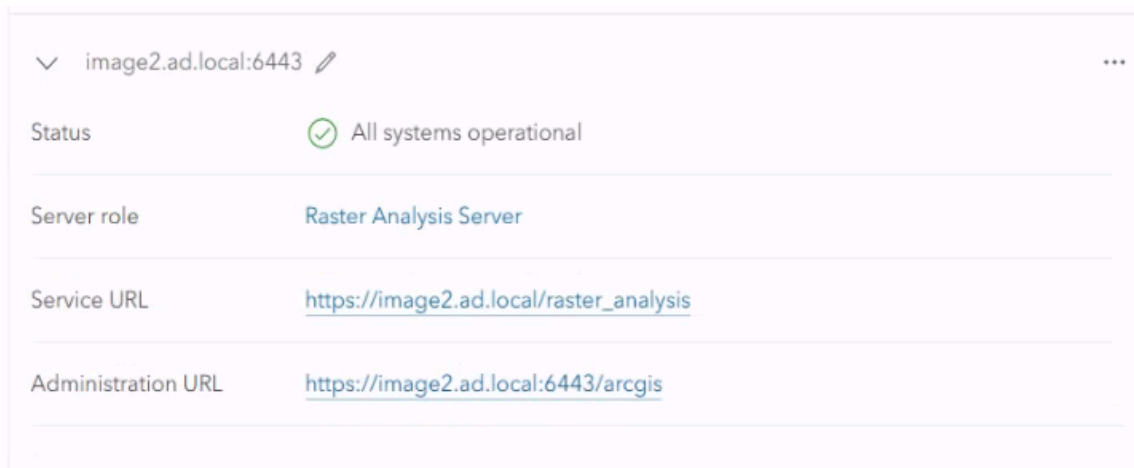
- i From the Configure Server Role tab, notice that the Raster Analysis Server role is assigned to Image.ad.local.

Before you can assign this role to Image2.ad.local, you will need to unassign the role from image.ad.local.

- j Click Done.
- k Under Image.ad.local, for Server Role, click Raster Analysis Server to edit the server role.
- l Disable the Raster Analysis Server role.
- m Verify that the Image Hosting Server role is still enabled.



- n Click Save.
- o Under Image2.ad.local, for Server Role, click Configure Server Role.
- p Enable the Raster Analysis Server role, and then click Save.



Now that your second ArcGIS Image Server site has been configured for the Raster Analysis

Server role, you can proceed.

- q Close the FederationURLs.txt file.

Step 3: Explore system raster tools

You will review the system services that are involved with raster capabilities in both Image Server machines. You will only enable those system raster services that are minimally required for each Image Server to perform the functions that it supports (image hosting for the IMAGE VM and raster analytics for the IMAGE2 VM).

- a Restore the IMAGE2 VM (lavender desktop).
- b In the web browser, refresh the ArcGIS Server Manager tab (https://image2.ad.local/raster_analysis/manager).
- c If you are unable to connect to ArcGIS Server Manager, clear the browser cache.
- d Sign in with the **portaladmin** username and **Esri.4.GIS** password.
- e From the Services tab, click the System folder and search for the **raster** services again.

2. What is the state of the system raster services now?

Because the IMAGE2 VM does not support a GPU and because raster hosting will be done on IMAGE, you do not need the last two services to be running.

- f For the RasterProcessingGPU and RasterRendering services, click the Stop button  .

The screenshot displays four service cards in the ArcGIS Server console:

- RasterAnalysisTools** (Geoprocessing Service): Status: Started. Instances Running: 0, Instances in Use: 0, Maximum Instances: 2.
- RasterProcessing** (Image Service): Status: Started. Instances Running: 0, Instances in Use: 0, Maximum Instances: 2.
- RasterProcessingGPU** (Image Service): Status: Stopped. Instances Running: 0, Instances in Use: 0, Maximum Instances: 0.
- RasterRendering** (Image Service): Status: Stopped. Instances Running: 0, Instances in Use: 0, Maximum Instances: 0.

Next, you will increase the number of instances of the RasterProcessing service. Instances allow ArcGIS Server to fulfill requests to services, so by increasing the number of instances, you increase the processing speed and capacity of raster analysis.

- g** For the RasterProcessing service, click the Edit Service button to view the properties of the RasterProcessing service.
- h** Click the Pooling tab, and then under Specify Number Of Instances, change Maximum Number Of Instances Per Machine to **4**.

The screenshot shows the configuration window for the RasterProcessing service, with the 'Pooling' tab selected. Under the 'Specify Number of Instances' section, the 'Maximum number of instances per machine' is set to 4.

| Specify Number of Instances | |
|--|--------------------------------|
| Minimum number of instances per machine: | <input type="text" value="0"/> |
| Maximum number of instances per machine: | <input type="text" value="4"/> |

- i** Click Save And Restart.

Now you will stop the unused system raster services on the image hosting server (IMAGE).

- j Open a new web browser tab and click the bookmark for IMAGE-Server Manager (<https://image.ad.local/image/manager>).
- k From the Services tab, click the System folder to view the system services on the IMAGE machine.
- l In the search field, type **raster** to expose the raster support services, as you did previously.
- m Stop all the services except for RasterRendering, which is required for the image hosting server.

The screenshot displays the Services console for ArcGIS Server, filtered for 'raster' services. Each service card includes an icon, a description, and a table of instance statistics.

| Service Name | Type | Status | Instances Running | Instances in Use | Maximum Instances |
|---------------------|-----------------------|---------|-------------------|------------------|-------------------|
| RasterAnalysisTools | Geoprocessing Service | Stopped | 0 | 0 | 0 |
| RasterProcessing | Image Service | Stopped | 0 | 0 | 0 |
| RasterProcessingGPU | Image Service | Stopped | 0 | 0 | 0 |
| RasterRendering | Image Service | Started | 0 | 0 | 2 |


- n Leave the web browser and web browser tabs open.

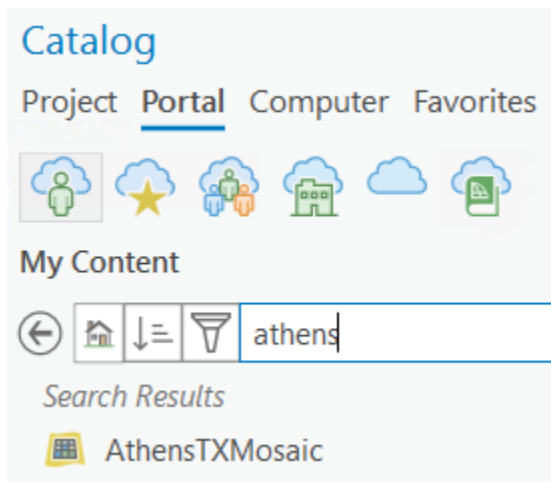
Step 4: Use Image Server to perform distributed raster analysis

Because you have federated an Image Server site (IMAGE2) and designated it as the raster analytics server for your ArcGIS Enterprise portal, users will have the ability to perform

distributed server-based processing on imagery and raster data. These capabilities are exposed in Map Viewer, but they are also available in ArcGIS Pro.

In this step, you will execute a distributed raster analysis job and then promptly move to the next step to observe the services in action using ArcGIS Server Manager.

- a Minimize the Remote Desktop Connection Manager window to view the host desktop.
- b On the host VM, start ArcGIS Pro.
- c In the top-right corner, click Sign In, and then sign in with the following credentials:
 - Username: **portaladmin**
 - Password: **Esri.4.GIS**
- d Under Recent Projects, click CourseProject.
- e In the Catalog pane, click the Portal tab.
- f Click the My Content button , if necessary.
- g In the Search My Content field, type **athens** and press Enter.



- h Add the AthensTXMosaic imagery layer to a new map.

Hint: Right-click the AthensTXMosaic imagery layer, point to Add To New, and choose Map.

- i From the Analysis tab, in the Portal group, click Raster Analysis, and then under Raster Analysis Tools, click Raster Functions.

The Raster Functions pane appears.

Raster functions apply processing directly to the pixels of imagery and raster datasets, as opposed to geoprocessing tools, which write out a new raster to disk.

- j** In the Raster Functions pane, expand Analysis, if necessary.
- k** Point to the NDVI Colorized function and read the description that appears.

You will use your imagery layer to create a Normalized Difference Vegetation Index (NDVI) raster. This raster will identify vegetation using two bands from your imagery: the red band (band 1) and the near-infrared band (band 4). The chlorophyll in healthy vegetation reflects more near-infrared light and absorbs more red light, so a ratio between these two bands is a common way to quantify vegetation from imagery.

- l** Click NDVI Colorized.
- m** In the Raster Functions pane, for the NDVI Colorized function, specify the following parameters:
 - Raster: AthensTXMosaic
 - Visible Band ID: 1
 - Infrared Band ID: 4
- n** Leave the Color Scheme Type and Colormap parameters as the default values.
- o** Under Output Layer Type, click the down arrow and read the descriptions for the available options.
- p** Choose Web Imagery Layer.

Raster

AthensTXMosaic

Visible Band ID

1

Infrared Band ID

4

Color Scheme Type

Colormap

Colormap

NDVI3

Output Layer Type

Web Imagery Layer

- q Click Next.
- r In the Raster Functions pane, specify the following parameters:
 - Name: **Athens_NDVI_Colorized**
 - Description: **Colorized NDVI analysis for Athens, Texas.**
 - Tags: **Vegetation, NDVI, Athens**
 - Sharing Options: ArcGIS Enterprise checked
- s Near the top of the Raster Functions pane, click the Settings tab.
- t Expand the Advanced Options section, if necessary.

You will now specify the number of raster processing service instances that will be used for processing this raster function. This parameter is an optional parameter, but you will set it to the

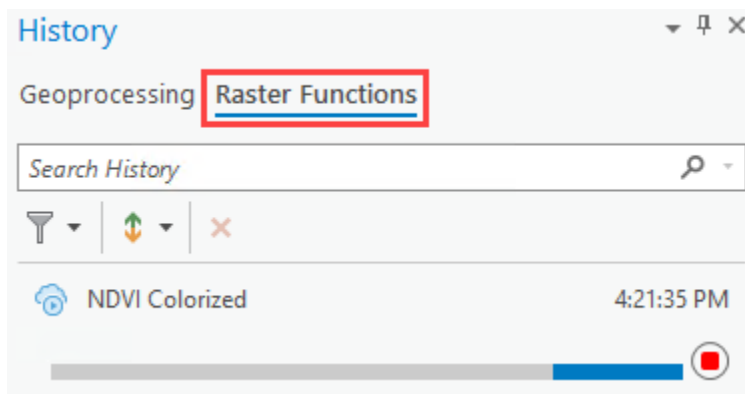
maximum number of instances (4) that you configured in a previous step to speed up the analysis.

- u For Parallel Processing Factor, type **4**.
- v Click Run.

i Submitted NDVI Colorized to portal for distributed raster analysis.

A notification confirms that the raster analysis job was submitted to the Image Server.

- w From the Analysis tab, in the Geoprocessing group, click History.
- x In the History pane, click the Raster Functions tab, as indicated in the following graphic, and confirm that the raster analysis job has been submitted.



- y Leave ArcGIS Pro open.

Because the job will begin being processed shortly, you will immediately inspect the raster analytics server.



Ensure that you perform the next step without any delay.

Step 5: Inspect running raster analysis services

You will now look at the activity in Server Manager caused by your analysis job.

- a Return to the IMAGE2 VM (lavender desktop).

- b Restore the web browser tab that shows the system raster services in the raster analytics server (ArcGIS Server Manager with the URL [https://image2.ad.local/raster_analysis/manager ...](https://image2.ad.local/raster_analysis/manager...)).
- c Search for the **raster** services again, if necessary.
- d Locate the RasterAnalysisTools service and the RasterProcessing service.

A single RasterAnalysisTools instance will start, divide the image extent, and begin to queue RasterProcessing tasks that will run in parallel.



You will need to wait a moment for all the instances to start and your screen to match the following graphic.

RasterAnalysisTools (Geoprocessing Service)

The RasterAnalysisTools service is used by ArcGIS Server to provide distributed raster analysis.

| | |
|--------------------|---------|
| Status: | Started |
| Instances Running: | 1 |
| Instances in Use: | 1 |
| Maximum Instances: | 2 |

👤 ▶ ■ ✕

RasterProcessing (Image Service)

The RasterProcessing service is used by ArcGIS Server to process raster data.

| | |
|--------------------|---------|
| Status: | Started |
| Instances Running: | 4 |
| Instances in Use: | 4 |
| Maximum Instances: | 4 |

👤 ▶ ■ ✕

- e When the number of running RasterProcessing tasks begins to reduce, restore the web browser tab for the image hosting server (ArcGIS Server Manager with the URL [https://image.ad.local/image/manager ...](https://image.ad.local/image/manager...)).
- f In the display of its system raster tools, watch the RasterRendering service until it indicates that an instance is running.

RasterRendering (Image Service)


The RasterRendering service is a multi-tenant raster rendering provider.

| | |
|--------------------|---------|
| Status: | Started |
| Instances Running: | 1 |
| Instances in Use: | 0 |
| Maximum Instances: | 2 |

👤 ▶ ■ ✕

When an instance of the RasterRendering service is running, the analytical results are ready to view.

- g Click the Hosted folder.
- h Confirm that the analysis result has a corresponding image service.




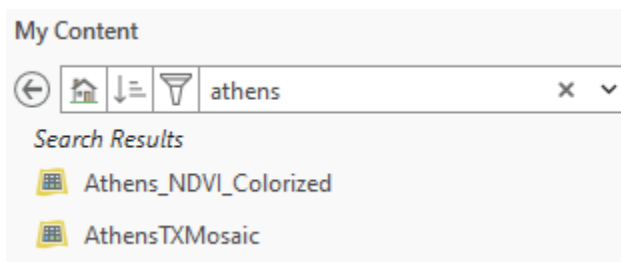
- i Close the web browser.

In this step, you observed the coordinated service activity on both of your ArcGIS Image Server sites.

Step 6: Explore imagery layer output


Finally, you will investigate the analysis result that was created by the NDVI Colorized raster function. This result is exposed as a web imagery layer, but the resulting data is stored in your raster store.

- a On the host VM, restore ArcGIS Pro.
- b In the Catalog pane, click the Portal tab, if necessary.
- c Click the Menu button  and choose Refresh.



The Athens_NDVI_Colorized imagery layer appears under My Content.

- d Right-click the Athens_NDVI_Colorized imagery layer and choose Add To Current Map.



The results of your analysis can now be accessed in the form of a web imagery layer that is shared with your ArcGIS organization.

- e Save your project, and then exit ArcGIS Pro.

In this exercise, you implemented workload separation between your two ArcGIS Image Server sites.

Observing a complex ArcGIS system

ArcGIS Enterprise systems are often complex; these systems may consist of multiple machines, include numerous federated ArcGIS Server sites, use high-availability design approaches, implement workload separation, and more. A complex ArcGIS Enterprise system can be cumbersome for administrators to consistently monitor, analyze, troubleshoot, and optimize. This notion is especially true if your organization needs to work with multiple stakeholders or use disparate monitoring and diagnostic tools across IT infrastructure.

ArcGIS Monitor is an application that supports the observability pillar of the ArcGIS Well-Architected Framework. This application allows administrators to obtain a holistic view of the health, performance, and usage of a complex ArcGIS Enterprise system through a single access point. ArcGIS Monitor can help you answer the following questions:

- Are the current system resources adequate to support load and still allow for optimal performance?
- What do the usage trends look like (over a day, a week, a month)?
- Are there any patterns that you can detect that point toward possible issues?
- Where are there bottlenecks that are causing high response times?

Observing a complex ArcGIS system (continued)

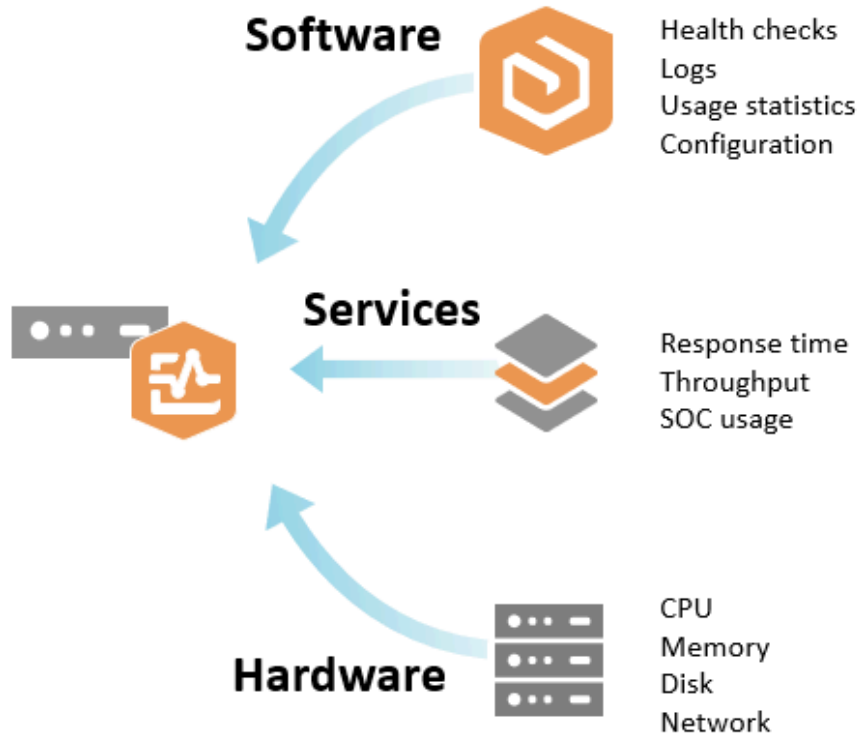


Figure 3.11. ArcGIS Monitor helps administrators identify and troubleshoot issues, optimize performance, and make informed decisions to enhance the efficiency and reliability of their ArcGIS system implementations.



ArcGIS Monitor - Esri Documentation: Introduction to ArcGIS Monitor

Observing a complex ArcGIS system video summary

ArcGIS Monitor is designed to collect, analyze, and optimize the health of your ArcGIS Enterprise components. These components include Portal for ArcGIS, ArcGIS Server, databases, and more. In real time, ArcGIS Monitor checks, analyzes, and reports on system health across your organization.

Analyzing system health

An administrator can create customizable views to analyze the health of select systems. The following views are among the possible views:

- **Alerts:** The number of alerts by type can be summarized for a desired time period. For example, you may want to show total alerts, critical alerts, warning alerts, and info alerts for the past 12 hours.
- **Information about a system:** Information about a given system can be summarized. For example, you may want to show how many members are part of the portal and summarize the number of members that are assigned to each user role and user type.
- **Resource utilization:** Charts summarizing CPU, memory utilized, and disk read/write used by the infrastructure can be included so that you can quickly view spikes in resource utilization.
- **Resource allocation:** The average number of ArcSOCs being used, minimum and maximum ArcGIS SOC instances by service, system response time, and service request rates are some of the metrics that can be included to aid decisions related to resource allocation.

Viewing reported alerts

When ArcGIS Monitor checks the system, any issues are reported as alerts. Administrators can view these alerts and filter the issues by type, such as critical, warning, or info. Each alert contains charts and metrics to better understand the history of the issue, which can help administrators address the problem before it impacts end users.

Lesson review

1. High availability is a design approach that allocates functions of your system to specific servers to prevent resource contention.
 - a. True
 - b. False

2. Which example is the most representative of workload separation?
 - a. Combining analytical and visual processing within the hosting server
 - b. Configuring GeoEvent and GeoAnalytics Servers together on the same machine
 - c. Federating a single ArcGIS Server site into an ArcGIS Enterprise deployment
 - d. Dedicating a faster CPU machine for analysis and a weaker machine for visualization

3. Which best practice is represented by an ArcGIS Enterprise deployment with one primary ArcGIS Data Store and one standby ArcGIS Data Store?
 - a. Workload separation
 - b. Load balancing
 - c. High availability
 - d. Capability delivery

Answers to lesson 3 questions

Explore ArcGIS Enterprise high availability (page 3-4)

1. What is the difference between high availability and disaster recovery?

High availability is focused on avoiding unplanned downtime. Disaster recovery is focused on retaining the data to restore a system after downtime.

2. What are two metrics that an organization might use to measure whether the system is meeting the necessary level of operational performance over time?

Availability targets and criticality tiers

3. What are six design constraints that can affect how high availability is achieved?

Business needs, deployment pattern, infrastructure, maintenance, software, and people and processes

Explore workload separation (page 3-13)

1. What are two common approaches to workload separation?

- **Isolate processes to prevent one process from overwhelming the other processes.**
- **Use server roles to enforce separation. For example, put imagery services on an ArcGIS Image Server rather than an ArcGIS GIS Server.**

2. Why might you want to isolate visualization and analysis tasks in your deployment?

Isolating analysis ensures that a large geoprocessing task will not affect mapping services, which could negatively impact other users.

Comparing high availability and workload separation (page 3-14)

| Characteristic | Distributed computing strategy |
|--|--------------------------------|
| Computationally intensive tasks cannot negatively impact other functions | Workload separation |

Table continued on next page

Answers to lesson 3 questions (continued)

| Characteristic | Distributed computing strategy |
|---|--|
| Eliminates a single point of failure through redundant infrastructure | High availability |
| Used to meet business-defined service level agreements (SLAs) | High availability and Workload separation |

Exercise 3: Implement workload separation (page 3-17)

1. What is the current state of the raster services?

Stopped

2. What is the state of the system raster services now?

Started

4

Web service fundamentals

Web services are the means by which ArcGIS Enterprise makes GIS resources available to client apps. Different types of services offer different capabilities that may be configured to meet your organization's business needs. In this lesson, you will explore how services interact with data and are exposed as web layers.

Topics covered

Web service fundamentals

User-managed and ArcGIS-managed data

Comparing services and web layers

Web service fundamentals

It is important to consider web services in relation to corresponding web layers in an Enterprise portal. When a user shares a web layer, a service is published and exposed through the web layer. When another user adds this layer to a map, the client app, such as Map Viewer, communicates directly with the underlying service. Although the web service provides access to the data, the web service itself is not a data storage location. Web services access data from a storage location, referred to in ArcGIS Enterprise as a data store. It is important to understand how web layers, web services, and data stores work together to power visualization, analytics, and storage in ArcGIS Enterprise.

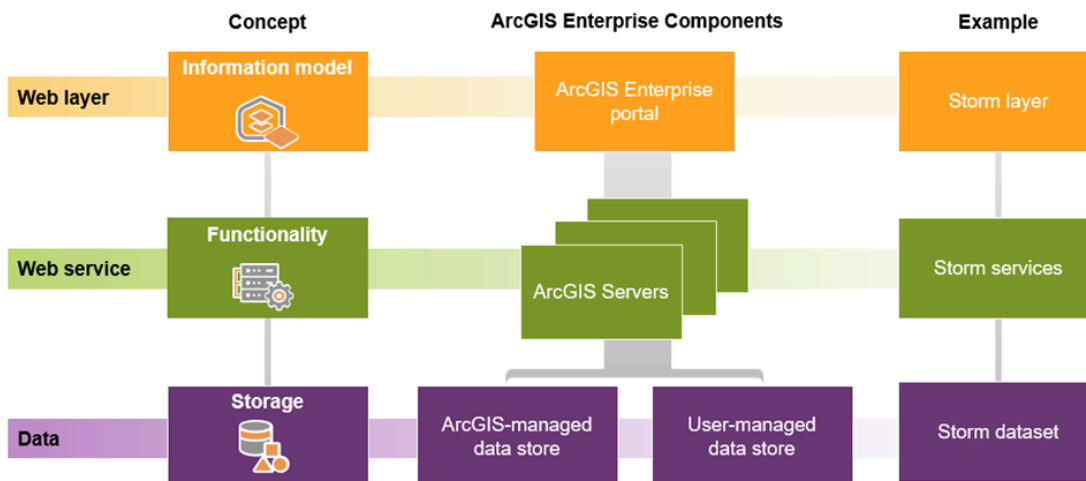


Figure 4.1. This graphic conceptualizes the hierarchy within which GIS services and layers exist, starting with the web layer and ending with the source data. In ArcGIS Enterprise, data stores provide data storage, ArcGIS Servers provide functionality through web services, and ArcGIS Enterprise portal exposes these resources in an information model as a web layer.

- Web services, powered by an ArcGIS Server site, take requests, fulfill the requests, and then send the results back to client applications, such as Map Viewer or ArcGIS Pro, that need to use the results.
- Web layers expose services as items in the Enterprise portal and are how portal users work with your data. Web layers contain metadata and tags, making them easily discoverable and simple to add to a map or an app.
- Data for web services are identified using data stores. Data stores contain information for

Web service fundamentals (continued)

ArcGIS Enterprise to find and access the data. There are two overarching ways in which data can be handled in ArcGIS Enterprise: user-managed and ArcGIS-managed.

Explore user-managed data and ArcGIS-managed data

It is important to consider business needs when determining where data will be stored and accessed using ArcGIS Enterprise. There are two ways in which data can be managed in ArcGIS Enterprise: user-managed and ArcGIS-managed. There are functionality differences that depend on whether the data should be user-managed or ArcGIS-managed.

User-managed

Data can be managed outside of ArcGIS Enterprise by an administrator, such as a SQL Server DBA. When you bring your own data, such as authoritative data in an enterprise geodatabase, to ArcGIS Enterprise, the data is user-managed.

ArcGIS-managed

Alternatively, data can be copied and managed in ArcGIS Enterprise. It forms the foundation of the ArcGIS Enterprise portal's self-service mapping and standard feature analysis model. If a workflow results in ArcGIS Enterprise creating a dataset, such as a portal analysis result, the data will be ArcGIS-managed.

You will review an Esri technical paper to learn about the differences between user-managed data and ArcGIS-managed data.

Instructions

- a** On the host VM, open File Explorer, and then browse to **C:\EsriTraining\EADM\OnlineFiles** and open the **DataInArcGIS.pdf** file.
- b** In the Esri technical paper, review pages 9 and 10.
- c** Use the information found on these pages to answer the following questions in your workbook.
- d** After you are finished, close the PDF file and close File Explorer.

Explore user-managed data and ArcGIS-managed data (continued)

1. When running analysis in ArcGIS Enterprise, which type of data will be the output?

2. What are some common reasons for storing data as user-managed?

3. Which type of data supports behaviors like topologies and versioning?



Esri technical paper: *Data in ArcGIS: User Managed and ArcGIS Managed*

How layers enable efficient exposure of services

There may be situations in which it is advantageous to create multiple layers that reference the same service.

Example of layers, a service, and data

A single service can expose multiple datasets, and multiple layers can reference a single service.

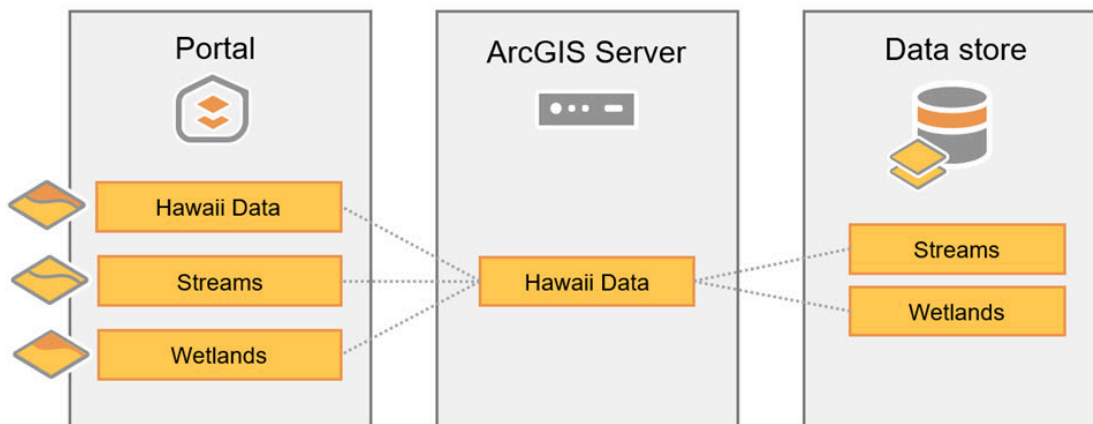


Figure 4.2. In this example, a data store may contain two datasets of differing geometries that are exposed in a single service. Any number of web layers that reference this service can then be created and are discoverable in a portal.

Using multiple layers to expose a service

The ability to use multiple layers to reference a single service allows you to publish fewer services. Web layers represent properties applied to data, such as the symbols and colors used, filters applied to the data, pop-up configurations, and more. For example, using the same underlying service, you can design multiple layers for use in different apps, with each layer showing a specific subset of the data with different symbology. Layers provide a way to customize visualization to meet the business need of the service.

You can identify which web service that a web layer is referencing from the web layer's item details in the portal. This action can be helpful in determining not only the origin of the underlying web service but also the underlying data represented in the web layer.

Understand the relationship between services and layers

Layers expose your services as items in a portal and are the building blocks of web maps and web scenes. Layers are used to interact with the data in your services, and the type of layer determines how you can interact with the layer's data.

You will examine your exercise environment to answer the following questions about the link between layers and services.

Instructions

- a** On the host VM, click the Remote Desktop Connection Manager icon and restore the EBASE VM (blue desktop).
- b** Open a web browser and click the ArcGIS Enterprise bookmark.
- c** Sign in with the **publisher** username and **Esri.4.GIS** password, if necessary.
- d** In ArcGIS Enterprise, click Content, and then locate and click the Cico Park -Trees Of Interest web map.
- e** For Cico Park -Trees Of Interest, click the Options button **⋮** and choose View Details.
- f** Under Layers, right-click Bald Cypress and choose Open Link In New Tab.
- g** Open the Oak link in a new tab.
- h** After you have finished answering the following questions, close the web browser.

Exposing services through layers

Perform the following steps to answer questions 1 and 2:

1. On each layer's item page, scroll down to the bottom of the page to find the URL to the corresponding web service.
2. To the right of URL, click View to navigate to the service that is listed in the ArcGIS REST Services Directory to answer the questions.



Understand the relationship between services and layers (continued)

1. **Can you have multiple layer items referencing the same service with different configurations? Why or why not?**
-

2. **Which type of service is exposed by a feature layer?**
-

Saving layer configurations in maps

Perform the following steps to answer question 3:

1. In the browser, restore the Cico Park tab, and then click Open In Map Viewer.
2. On the right, from the Settings toolbar, click the Styles button .
3. Under Pick A Style, click Style Options, edit the symbol style, and then click Done three times.
4. On the left, from the Contents toolbar, click the Save And Open button and choose Save As.
5. Leave all defaults and click Save to create a copy of the map.
6. In the top-left corner, click the Menu button  and choose Content.
7. Reopen the original Cico Park -Trees Of Interest web map with Map Viewer.

3. **Are the original layers modified when you save the map?**
-
-

Exercise 4

Compare layers, services, and data

From an end-user perspective, layers are easy to discover and use, regardless of their type. Different types of layers can be added to a map, used to perform analysis, and even shared without in-depth knowledge of the underlying service type and data source.

As an administrator, however, it is important to understand the relationship between layers, web services, and data sources. Different types of services have different properties that affect the user experience as well as system resource utilization.

In this exercise, you will perform the following tasks:

- Publish feature layers that use ArcGIS-managed and user-managed data.
- Investigate feature layers that use ArcGIS-managed and user-managed data.

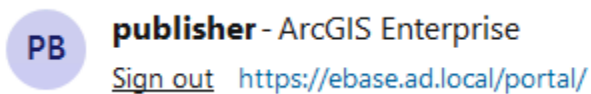
Step 1: Publish a feature layer that uses ArcGIS-managed data

To begin, you will publish a feature layer from a map layer in ArcGIS Pro. The source of the feature layer is a shapefile, and the result will be a feature layer that uses ArcGIS-managed data.

- a On the host VM, start ArcGIS Pro.

You will sign in with the publisher user.

- b In the top-right corner, click Sign Out, if necessary.
- c Click Sign In, and then sign in with the following credentials:
 - Username: **publisher**
 - Password: **Esri.4.GIS**



- d Under Recent Projects, click CourseProject.
- e Close any open maps in the project.

Within this project, you will find a map titled Houston Public Services. This map contains a Libraries layer and a Metro Bus Routes layer. The Libraries layer is sourced from a shapefile, and the Metro Bus Routes layer is sourced from an Enterprise geodatabase.

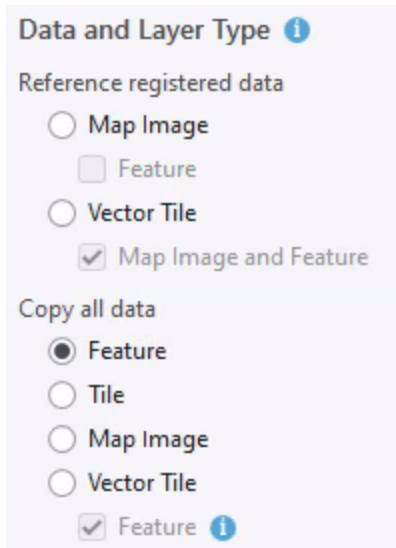
- f In the Catalog pane, on the Project tab, expand Maps.
- g Right-click the Houston Public Services map and choose Open.



- h** In the Contents pane, right-click the Libraries layer and choose Properties.
- i** In the Layer Properties dialog box, click the Source tab.
- j** Expand the Layer Properties dialog box, if necessary, to see the full file path of the shapefile.
- k** Review the information in the Data Source section.

This layer is stored locally as a shapefile. When you publish this data as a web layer, you will choose to copy the data to the server, rather than referencing the data.

- l** Click Cancel to close the Layer Properties dialog box.
- m** In the Contents pane, right-click the Libraries layer, point to Sharing, and choose Share As Web Layer.
- n** In the Share As Web Layer pane, under Data And Layer Type, in the Copy All Data group, select the Feature option.



The Libraries layer will be copied and managed by ArcGIS Enterprise, and the feature layer will use ArcGIS-managed data.

- Under Location, for Portal Folder, type **FeatureLayers**.

This parameter will create a new folder in your Enterprise portal.

- Click Analyze to ensure that your data is ready to be published.

- Click Publish, and then wait for the process to complete before continuing.

In this step, you shared a feature layer to the portal, choosing to Copy All Data. This action will result in a feature layer that uses ArcGIS-managed data. The data being copied into ArcGIS Enterprise will be stored in the relational ArcGIS Data Store.

Step 2: Investigate a feature layer that uses ArcGIS-managed data

Now that you have published a feature layer that uses ArcGIS-managed data, you will investigate the results in ArcGIS Enterprise portal and ArcGIS Server Manager.

- Restore the EBASE VM (blue desktop).


The feature layers that you previously shared were published to the base deployment, so you will investigate the layers and services through the EBASE VM.

First, you will investigate the feature layer exposed in ArcGIS Enterprise.

- b Open a web browser and, if necessary, click the ArcGIS Enterprise bookmark (<https://ebase.ad.local/portal/home>).

You will sign in with the publisher credentials.

- c Sign out of portaladmin, if necessary.
- d Click Sign In, and then sign in with the **publisher** username and **Esri.4.GIS** password.
- e Click Content, and then under Folders, click the FeatureLayers folder.
- f Notice the Libraries feature layer and service definition items.



Within the FeatureLayers folder, there are a pair of items with the same title. When you published and chose to copy the data to your server, you created the feature layer that uses ArcGIS-managed data and service definition items titled Libraries.

- g Click the Libraries feature layer item to view the item page.
- h On the right side of the item page, review the series of buttons that expose different capabilities of this feature layer.

You will notice that there are several options made available on the item page, including the ability to publish additional layers, create view layers, and export your data to a variety of different formats.

- i Click the Settings tab to review the configurable settings that are available for the feature layer.

The four notable settings that are available through your Enterprise portal are the abilities to enable editing, track changes to the data, enable sync for offline use and collaborations, and allow others to export data.

Now you will investigate the Libraries feature layer's web service.


- j Open a new web browser tab and click the EBASE-Server Manager bookmark to go to ArcGIS Server Manager (<https://ebase.ad.local/server/manager>).

A banner may appear that warns you that you cannot alter any server settings because you are not signed in as a portal administrator. This message is a consequence of being signed in as the publisher in the same browser session. You do not need to alter any settings, so you will continue as a publisher.

- k Under Folders, click the Hosted folder.

The Libraries feature layer and service uses ArcGIS Managed-data and is considered hosted. By default, all hosted layer services will be housed in the Hosted folder.

- l Click Libraries to view the service details.
- m Review the available options for the Libraries service.



There are no configurable options available through ArcGIS Server Manager. This fact is highlighted by the grayed-out Save And Restart button because there are no settings to modify.

Step 3: Publish a feature layer with user-managed data

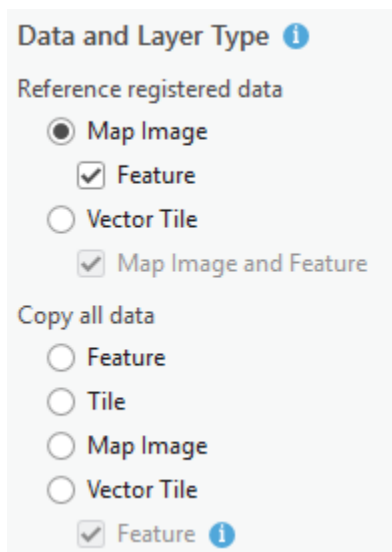
Next, you will publish a second feature layer from a map layer in ArcGIS Pro. The source of this map layer will be a feature class in an enterprise geodatabase.

- a On the host VM, restore ArcGIS Pro with the CourseProject project open.
- b In the Contents pane, right-click the Metro Bus Routes layer and choose Properties.
- c In the Layer Properties dialog box, click the Source tab, if necessary.
- d Review the information in the Data Source section.

This layer is stored within an enterprise geodatabase. When you publish this data as a web layer, you will choose to reference this registered data, rather than copying the data to the server. In addition to the location of the source data, you will notice that the enterprise geodatabase also

provides information such as privileges that are configured for this feature class.

- e Click Cancel to close the Layer Properties dialog box.
- f In the Contents pane, right-click the Metro Bus Routes layer, point to Sharing, and choose Share As Web Layer.
- g In the Share As Web Layer pane, under Data And Layer Type, in the Reference Registered Data group, check the box for the Feature option.



You will notice that unlike the Feature option in the Copy All Data group, the Feature option in the Reference Registered Data group requires you to select Map Image in addition to Feature. When sharing a feature layer that references registered data to ArcGIS Enterprise, it is shared as a dependent layer to a map image layer.

- h Under Location, for Portal Folder, choose the folder that you previously created named FeatureLayers.
- i Under Sharing Level, select Organization.

You are sharing the map image layer and feature layer with the organization. These layers will be used in future analysis and, therefore, will need to be shared with the ArcGIS Enterprise organization or with the appropriate group for access.

- j Click Analyze to ensure that your data is ready to be published.
- k Click Publish, and then wait for the process to complete before continuing.





- 1 After the layer has been published, save the project, if necessary, and then exit ArcGIS Pro.

In this step, you shared a map image layer and a feature layer to the portal, choosing to Reference Registered Data. This action will result in a map image layer and feature layer using user-managed data, with the data remaining in your Enterprise geodatabase.

Step 4: Investigate a feature layer that uses user-managed data

Next, you will investigate a feature layer that references data in an enterprise geodatabase.

- a Restore the EBASE VM (blue desktop) and open the ArcGIS Enterprise web browser tab.
- b Click Content, and then under Folders, click the FeatureLayers folder, if necessary.
- c Notice the Metro Bus Routes feature layer and map image layer items.

| Selected (0) 1-4 of 4 | |
|---|--|
| Title | |
| <input type="checkbox"/> Metro Bus Routes |  Feature Layer |
| <input type="checkbox"/> Metro Bus Routes |  Map Image Layer |
| <input type="checkbox"/> Libraries |  Feature Layer (hosted) |
| <input type="checkbox"/> Libraries |  Service Definition |

When you published the layer, you created the feature layer and map image layer items titled Metro Bus Routes.

- d Click the Metro Bus Routes feature layer to view the item page.

You will notice that there are fewer options available for this non-hosted feature layer compared to those that are available with the Libraries hosted feature layer.

- e Click the Settings tab.
- f Under Feature Layer, expand Configuration.

Like with a hosted feature layer, you can enable editing, enable sync for offline use and collaborations, and allow others to export data.

You will now investigate the web service for the Metro Bus Routes feature layer. When you published Metro Bus Routes, you selected the option to publish a feature layer with the map service option. Therefore, the feature service is managed through the map service in ArcGIS Server Manager.

- g** On the ArcGIS Server Manager web browser tab, within the Site (Root) folder, locate the Metro_Bus_Routes map service.
- h** Click the Metro_Bus_Routes map service and review the available options.

Unlike with the hosted feature service, several configurable options are available through ArcGIS Server Manager.

- i** Click Capabilities.

You will notice that mapping and feature access are enabled for the Metro_Bus_Routes map service. The mapping option corresponds to the map image layer in the portal, and the feature access option corresponds to the feature layer.

- j** On the EBASE VM (blue desktop), close the web browser.

In this exercise, you published and reviewed available configuration settings for ArcGIS-managed and user-managed layers and services.

Lesson review

1. True or false: When you share a new web layer from ArcGIS Pro, a new web service is published and exposed as a web layer item in the Enterprise portal.
 - a. True
 - b. False

2. Which ArcGIS software component exposes data as services?
 - a. ArcGIS Data Store
 - b. Portal for ArcGIS
 - c. ArcGIS Server
 - d. ArcGIS Web Adaptor

3. Which layer utilizes user-managed data?
 - a. A feature layer published with Copy All Data from ArcGIS Pro
 - b. A map image layer referencing an enterprise geodatabase
 - c. A feature layer that was created using a CSV file in the portal
 - d. An imagery layer output from raster analysis tools in the portal

Answers to lesson 4 questions

Explore user-managed data and ArcGIS-managed data (page 4-4)

1. When running analysis in ArcGIS Enterprise, which type of data will be the output?

ArcGIS-managed, specifically hosted, layers; all analysis conducted in ArcGIS Enterprise will create an output that is stored and managed by ArcGIS Enterprise.

2. What are some common reasons for storing data as user-managed?

Answers may vary, but common reasons include:

- **Need for versioning, topology, archiving, and other advanced geodatabase capabilities**
- **Need for direct SQL access**
- **Need to use server-side rendering (for example, to use the Maplex Label Engine to provide precise control over label placement in a map service)**
- **Workflows that include users collaborating on a single set of data**

3. Which type of data supports behaviors like topologies and versioning?

User-managed data (enterprise geodatabase) supports these behaviors.

Understand the relationship between services and layers (page 4-7)

1. Can you have multiple layer items referencing the same service with different configurations? Why or why not?

Yes; the settings are stored with the layer instead of the service.

2. Which type of service is exposed by a feature layer?

A feature service is exposed by a feature layer.

3. Are the original layers modified when you save the map?

No; the original layers are not modified because any changes to the layers' symbology are stored with the new web map.

5

User-managed data

As you learned previously, there are two overarching ways in which data can be handled in ArcGIS Enterprise: user-managed and ArcGIS-managed. Both data storage types are essential to the implementation of a robust, modern, and web services-oriented architecture.

The choice of user-managed data will influence the publishing, management, and capability options of your web layers and services; therefore, it is imperative to understand the impacts. This lesson will explore types of user-managed data, how to configure the data to make it accessible by ArcGIS Enterprise, and how to configure services that reference user-managed data to achieve your business needs.

Topics covered

Types of user-managed data

Referencing registered data

Branch versioning

Types of user-managed data

As you previously learned, user-managed data can be leveraged in cases where a particular business need justifies its use, such as direct SQL access to a relational database. User-managed data sources typically require more work to maintain, as administrators provide and manage the underlying storage infrastructure. There are several types of user-managed data stores that allow you to integrate your user-managed data into ArcGIS Enterprise.




Figure 5.1. User-managed data can be stored in databases, folders, cloud stores, big data file shares, and NoSQL databases.

The following table provides information for each type of user-managed data store.

Types of user-managed data (continued)

| Data store type | Data source | Notes |
|---------------------|--|--|
| Database | Relational databases such as SQL Server, Oracle, and PostgreSQL | <ul style="list-style-type: none"> • Can contain enterprise geodatabases, which are ArcGIS technology built on top of relational database systems • Database does NOT include file geodatabases • Only way to publish a feature layer that uses user-managed data <p>Note: Provides capabilities such as versioning and topologies</p> |
| Folder | Registered file share location | <ul style="list-style-type: none"> • Can contain various data types such as shapefiles, rasters, and file geodatabases |
| Cloud store | Cloud-based storage location such as an Amazon S3 bucket or Microsoft Azure Blob storage container | <ul style="list-style-type: none"> • Can be used as a raster store, big data file share, or to store map and image caches |
| Big data file share | A folder, cloud store, HDFS location, Hive | <ul style="list-style-type: none"> • Store collections of delimited files (for example, .csv and .tsv), ORC, parquet, and shapefiles • ArcGIS GeoAnalytics Server can read from big data file shares and optionally write output to big data file shares |

Table continued on next page

Types of user-managed data (continued)

| Data store type | Data source | Notes |
|-----------------|----------------|--|
| NoSQL | Neo4j database | <ul style="list-style-type: none">• ArcGIS Knowledge Server can read user-managed data from an existing Neo4j database to create a knowledge graph |

Which user-managed data sources are most useful for your organization?

Referencing registered data from ArcGIS Pro

For a service to use user-managed data, the data needs to be registered with ArcGIS Enterprise. To share a layer that references registered data, you must first register the source as a data store with the ArcGIS Server site that you intend to publish the service to. Registering your data sources provides ArcGIS Server with a list of locations to access to populate web services with data. You can register a database, a folder, a cloud store, big data file share, or NoSQL database with ArcGIS Server.

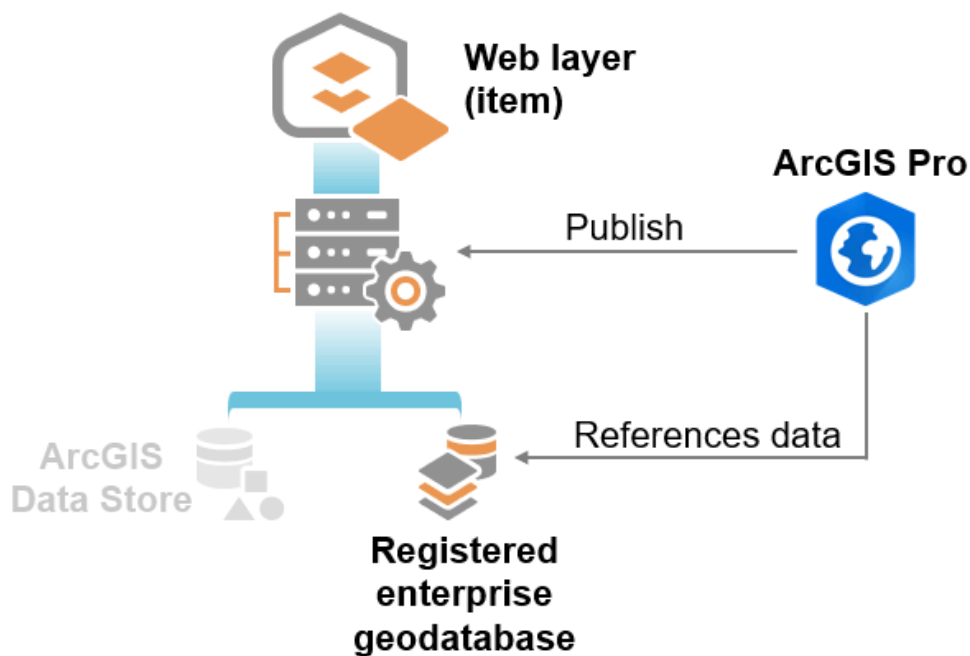


Figure 5.2. ArcGIS Pro can be used to publish services that reference a registered data store to ArcGIS Server.

A registered data store must be in a location that is accessible to ArcGIS Server, such as a shared

Referencing registered data from ArcGIS Pro (continued)

directory. Be sure to grant the ArcGIS Server account access to the data; for example, if your data is stored in a folder that you access using operating system authentication, you must grant the ArcGIS Server account permissions to this folder.



ArcGIS Enterprise Help: *Make your data accessible to ArcGIS Server*

Capabilities for services that use user-managed data

The capabilities that you provide to end users through web services is strongly driven by your organization's business needs. There are various properties that can be configured for a service using user-managed data that impact what a user can or cannot do with the service. Below are some of the capabilities that can be configured for the service:

- Editing/Feature Access
- Export and download
- The ability to change symbology
- Open Geospatial Consortium (OGC) protocols, such as WMS and WCS
- Caching

Publishing a service that references user-managed data impacts which service properties can be configured for the service. As the item owner or administrator, you may need to modify the properties of these services before and after the services have been published. The properties for a service using user-managed data will be modified during publishing in ArcGIS Pro or modified after publishing in ArcGIS Server Manager.

Versioning

Like other advanced enterprise geodatabase capabilities, a service with versioning capabilities is only available if the data is user-managed. Versioning requires data to be stored in an enterprise geodatabase, and it should be registered as a user-managed data store. Versioning is a framework in which multiple editors can work simultaneously on a single feature class in a highly isolated fashion without creating copies of the data.



An Advanced Editing user type extension is required to create and edit named versions.

Types of versioning

ArcGIS supports two types of versioning models: traditional versioning and branch versioning. Traditional versioning allows operations to be performed at the data level (within a multiuser geodatabase/RDBMS). However, branch versioning enables operations through feature services instead of directly on the RDBMS level, facilitating a Web GIS paradigm. Because branch versioning is designed for services-based architecture, it is inherently more integrated into ArcGIS Enterprise; therefore, it will be the versioning type that you will use in this lesson.

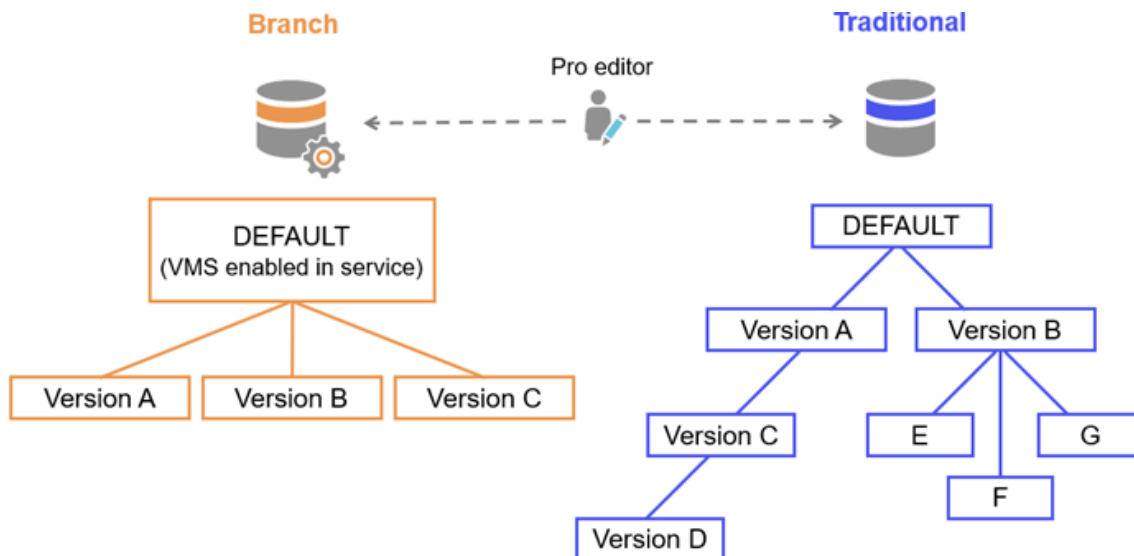


Figure 5.3. This example illustrates the difference between traditional versioning and branch versioning. While the editing experience is nearly identical, branch versioning provides a simple architecture that operates using feature services.

Versioning (continued)

Branch versioning

For branch versioning, upon publishing, the service property will need to be configured to take advantage of the versioning capability. Branch versioning has several defining characteristics and related administrative considerations.

Branch versioning

| Characteristic | Administrative considerations |
|--|--|
| Operates on feature services | <ul style="list-style-type: none"> • No direct database connection is required • Key version administration tasks are performed through services, enabled by the Version Management capability (VMS) |
| Simple table structure (no adds, deletes, or state tables) | <ul style="list-style-type: none"> • Does not need to be compressed • No performance degradation over time due to complex version tree structures and sets of tables |
| Child versions are derived only from DEFAULT | <ul style="list-style-type: none"> • Potentially simpler architecture of version tree • Editors work with their own version and post their changes to DEFAULT |

Branch versioning may have key differences from traditional versioning, but much of the editor's experience with branch versioning is identical to traditional editing workflows. In branch versioning, editors make edits, reconcile, and (optionally) post with DEFAULT, just as they would do with traditional versions.

Configure branch versioning

There are several steps that must be performed in ArcGIS Pro to ensure that the data is ready to share as a branch-versioned feature layer.

To configure branch versioning in ArcGIS Pro, you must perform the following steps:

Versioning (continued)

1. Register data residing in the enterprise geodatabase as a user-managed data store.
2. Make a connection to the enterprise geodatabase as the dataset owner.
3. Register the dataset as branch versioned, which adds a Global ID field to the dataset and enables editor tracking and archiving.
4. Share branch-versioned data to your Enterprise portal as a feature service that references registered data and set the service property to version management.
5. Use the web feature layer to perform branch version workflows.




Figure 5.4. Branch versioning requires the data to be user-managed; therefore, the first step requires the data to be registered with ArcGIS Enterprise. Next, ArcGIS Pro is used to make a database connection using the branch versioning type. Then, the dataset needs to be registered as branch versioned, which also adds a Global ID field to the dataset and enables editor tracking and archiving. The dataset can then be shared as a web layer from ArcGIS Pro and subsequently used in versioned editing workflows.

 **ArcGIS Pro Help: Branch version scenarios**

Exercise 5

Perform branch versioning workflows


Your organization wants to configure branch versioning for multiple editors to simultaneously update city boundary data, as necessary. To enable branch versioning, you know that the first step is to ensure that the city boundary data is accessible to ArcGIS Enterprise. You begin the process by registering the enterprise geodatabase as a user-managed data store. Now that the underlying data has already been registered with ArcGIS Enterprise, you will begin the exercise by performing the steps necessary to register the dataset as branch versioned before sharing it as a web layer to portal. When the feature service is published, you will create a named version. By the end of the exercise, you will have service configured for branch versioning and ready for versioned editing workflows.


In this exercise, you will perform the following tasks:

- Share data as branch versioned.
- Manage branch versions.

Step 1: Register a dataset as branch versioned

ArcGIS Enterprise already has access to the city boundary data. The enterprise geodatabase that underlying data resides in has been registered as a user-managed data store. To begin, you will set up the data to take advantage of branch versioning by registering the data as branch versioned.

- a On the host VM, start ArcGIS Pro.
 - b Sign in with the **publisher** username and **Esri.4.GIS** password, if necessary.
-  You may need to sign out from the portaladmin account before signing in as publisher.
- c Open CourseProject and close all open map views.
 - d In the Catalog pane, expand Databases and locate the Publisher@EDATA.sde connection.



- e Right-click Publisher@EDATA.sde and choose Connection Properties.

1. What is the database platform for this enterprise geodatabase?

In the Database Connection dialog box, you can see that you are connected to the edata enterprise geodatabase as the database user named publisher (note that this user is different from the Enterprise portal member named publisher). The publisher database user will be the data owner of the feature class that you load into this enterprise geodatabase.

- f Click Cancel to close the Database Connection dialog box.
- g In the Catalog pane, expand the Publisher@EDATA.sde connection, and then right-click the Edata.publisher.Cityboundary feature class and choose Manage.

The Feature Class Properties dialog box opens, displaying the Manage tab.

- h Under Manage Geodatabase Functionality, check the box for Versioning and confirm that the Branch option is selected.

Manage geodatabase functionality

- Versioning
- Branch i
- Traditional

2. Which additional options are automatically enabled?

- i Click OK.


You will see several green check marks appear as the dataset is configured to support branch versioning, and the Feature Class Properties dialog box will close when the process is complete.

- j In the Catalog pane, right-click the Publisher@EDATA.sde connection and choose Geodatabase Connection Properties.



If Geodatabase Connection Properties is grayed out, try saving the project.

- k In the Geodatabase Connection Properties window, for Versioning Type, click Branch, if necessary.



The geodatabase workspace must be set to Branch to share the branch-versioned layer to your Enterprise portal.

- l Click OK to close the Geodatabase Connection Properties window.
- m Add the Cityboundary layer to a new map.

Hint: Right-click the Edata.publisher.Cityboundary feature class, point to Add To New, and choose Map.

The Cityboundary feature class is now registered as a branch-versioned dataset, and the geodatabase versioning type is set to branch.

Another requirement for publishing is to configure the map to assign unique numeric IDs.

- n** On the Contents pane, right-click the name of the map and choose Properties.
- o** On the General tab, name the map **BranchVersioning**.
- p** Check the box for Allow Assignment Of Unique Numeric IDs For Sharing Web Layers, and then click OK.
- q** Save the project.

You are ready to share the data to the Enterprise portal and configure the service property.

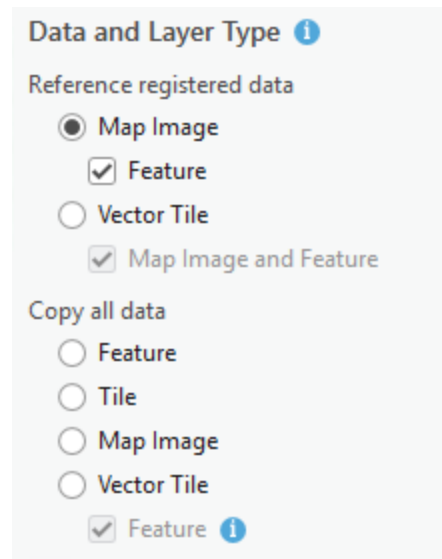
Step 2: Share a branch-versioned layer

Now that the feature class is configured to support branch versioning, you must share the feature class as a feature layer, which will be used in the editing process. Branch versioning differs from traditional versioning in that branch-versioned data is accessed through a feature layer and its underlying service. Therefore, you will share the data to be edited with the Enterprise portal (and by extension, the hosting server).

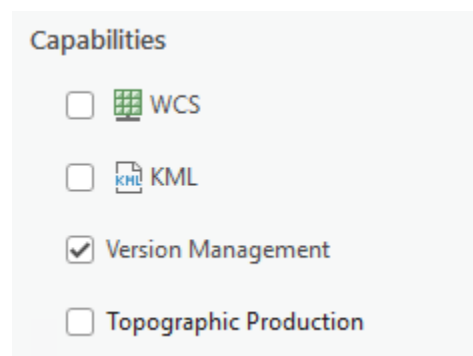
- a** In the Contents pane, click the symbol for the Cityboundary layer to open the Symbology pane.
- b** In the Symbology pane, click the Properties tab, and then for Color, choose any shade of blue.
- c** For Line Width, set it to 3 pt, and then click Apply.
- d** Close the Symbology pane.

Your layer is now ready to be shared to your Enterprise portal.

- e** In the Contents pane, right-click the Cityboundary layer, point to Sharing, and choose Share As Web Layer.
- f** In the Share As Web Layer pane, in the Data And Layer Type section, under Reference Registered Data, verify that Map Image is selected, and then check the box for Feature.




- g** Under Sharing Level, select Organization.
- h** Accept the remaining default settings but do not click Publish yet.
- i** Near the top of the Share As Web Layer pane, click the Configuration tab, and then under Capabilities, check the box for Version Management, if necessary.



If the Version Management option does not appear, try re-adding the layer to the map.

- j** From the Configuration tab, click the Configure Pooling button , as shown in the following graphic.



- k Under Instance Type, select Dedicated Instance, if necessary.

The shared instance pool does not support the Version Management capability. You will learn about the shared instance pool and the difference between dedicated and shared instances in a subsequent lesson.

- l Click Analyze and notice the error about registering the data.
- m Right-click the error and choose Register Data Source With Server.
- n In the Add Data Store dialog box, for Title, type **Manhattan**, and then for Tags, type **branch versioning**.
- o Verify that the box next to the hosting server is checked, and then click Validate.

Add data store [X]

Provide connection details for your data store

Title:

Tags:

Portal Folder: [Folder icon]

Publisher database connection

Server database connection

Same as publisher database connection
 New connection

Select the servers to which you want to add your data store

| <input type="checkbox"/> | Server | Status |
|-------------------------------------|--|--------|
| <input checked="" type="checkbox"/> | https://ebase.ad.local/server (Hosting Server) | ✓ |



If the validation takes longer than a couple minutes, you can proceed with the exercise. This action will not affect your ability to register the data source.

- p Click Create.

You are now ready to publish the service.

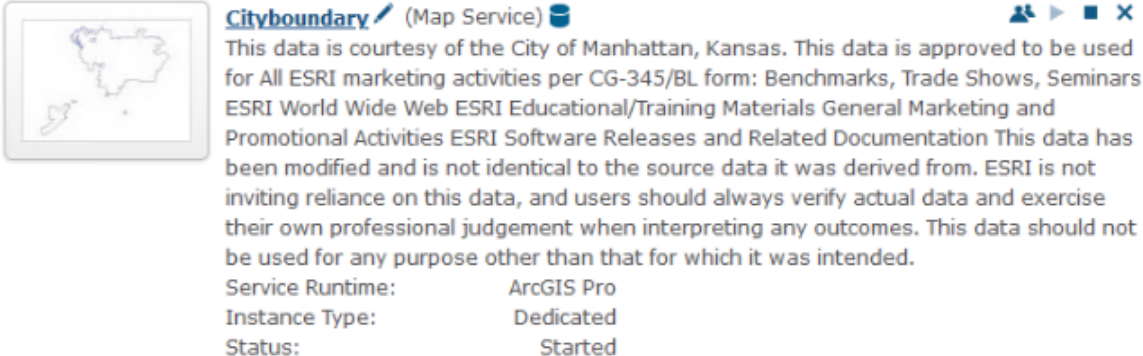
- q In the Share As Web Layer pane, click Publish.

In this step, you shared a web layer to your Enterprise portal that can support branch-versioned workflows.

Step 3: Inspect branch management capabilities in ArcGIS Server Manager

Next, you will explore the capabilities of the service that you published using ArcGIS Server Manager.

- a After the data has completed the publishing process, on the EBASE VM (blue desktop), open a web browser, and then click the EBASE-Server Manager bookmark.
- b Sign in to the portal as **publisher** with the **Esri.4.GIS** password.
- c From the Services tab, locate the Cityboundary map service.



Cityboundary (Map Service)

This data is courtesy of the City of Manhattan, Kansas. This data is approved to be used for All ESRI marketing activities per CG-345/BL form: Benchmarks, Trade Shows, Seminars ESRI World Wide Web ESRI Educational/Training Materials General Marketing and Promotional Activities ESRI Software Releases and Related Documentation This data has been modified and is not identical to the source data it was derived from. ESRI is not inviting reliance on this data, and users should always verify actual data and exercise their own professional judgement when interpreting any outcomes. This data should not be used for any purpose other than that for which it was intended.

| | |
|------------------|------------|
| Service Runtime: | ArcGIS Pro |
| Instance Type: | Dedicated |
| Status: | Started |

This service references the user-managed data that resides in your enterprise geodatabase.

- d Click the Cityboundary link, and then click the Capabilities tab.
- e Under Select And Configure Capabilities, click Version Management to select it.

Select and configure capabilities

- | | |
|--|--|
| <input checked="" type="checkbox"/> Mapping (always enabled) | <input type="checkbox"/> Validation |
| <input type="checkbox"/> WCS | <input type="checkbox"/> WFS |
| <input type="checkbox"/> Network Analysis | <input type="checkbox"/> Topographic Production |
| <input checked="" type="checkbox"/> Version Management | <input type="checkbox"/> Network Diagrams |
| <input type="checkbox"/> Trace Network | <input type="checkbox"/> KML |
| <input type="checkbox"/> Linear Referencing | <input type="checkbox"/> WMS |
| <input type="checkbox"/> OGC Features | <input type="checkbox"/> Utility Network |
| <input type="checkbox"/> Parcel Fabric | <input checked="" type="checkbox"/> Feature Access |

Version Management Configuration

URLs

REST URL: <https://ebase.ad.local/server/rest/services/Cityboundary/VersionManagementServer>

- f** Under Version Management Configuration, click the link next to REST URL to view the REST endpoint for the version management service.
- g** Identify the supported operations for the version management service.

This version management service allows branch versions to be managed through a services client, such as ArcGIS Pro.

Next, you will explore the feature server item.

- h** Near the top of the page, click the Services link, and then under Services, click the Cityboundary (FeatureServer) link.

ArcGIS REST Services Directory Logged in user : publisher | [Logout](#) |

[Home](#) > [services](#) > [Cityboundary \(FeatureServer\)](#) [API Reference](#)

[JSON](#) | [SOAP](#)

Cityboundary (FeatureServer)

View In: [ArcGIS JavaScript](#) [ArcGIS Online Map Viewer](#)

Service Description: This dataset represents the city boundary for the City of Manhattan, Kansas. The data and related materials are made available through Esri (<http://www.esri.com>) and are intended for educational purposes only (see Access and use limitations section).


- i Under the Layers section, click the link for Cityboundary, and then scroll to the bottom of the page.
- j Click Update Features and notice the Geodatabase Version Name box, as indicated in the following graphic.

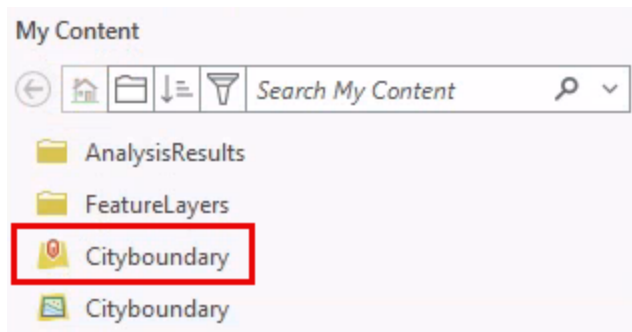
The screenshot shows the 'Update Features' dialog box. It has a 'Features:' section with a large empty text area. Below that is the 'Geodatabase Version Name:' field, which is highlighted with a red rectangle. Underneath are several options: 'Rollback on Failure' with radio buttons for 'True' (selected) and 'False'; 'Time Reference UnknownClient' with radio buttons for 'True' and 'False' (selected); 'True Curve Client' with radio buttons for 'True' (selected) and 'False'; and a 'Format:' dropdown menu set to 'HTML'. At the bottom left is an 'Update Features' button.

Because branch versioning is enabled, a specific version can be specified for a given edit.

Step 4: Create versions

Now you will learn how to control which type of versioned workflow that you will be conducting in ArcGIS Pro. Because branch versioning operations are performed through feature services, you will remove the source data from the map and add the feature layer.

- a On the host VM, restore ArcGIS Pro.
- b Close the Share As Web Layer pane.
- c In the Contents pane, click the List By Data Source button .
- d Right-click the Cityboundary layer and choose Remove.
- e In the Catalog pane, click the Portal tab, and then locate the Cityboundary feature layer.




Hint: If necessary, in the search field, type **city** and press Enter.

- f** Drag the Cityboundary feature layer into the map.

You will notice that the symbology is set properly, based on how you set it when sharing the data.

- g** In the Contents pane, select the sde.DEFAULT (Cityboundary) connection.




The Versioning tab appears at the top of ArcGIS Pro.

- h** From the Versioning tab, in the Versioning group, click New Version.



Because ArcGIS Pro is context sensitive, the Versioning tab will only appear when the connection is selected in the Contents pane.

- i** In the New Version dialog box, for Name, type **Version1**.
- j** For Description, type **My first branch version**.
- k** For Access Permission, verify that it is set to Private.
- l** Check the box for Change To This New Version.
- m** Click OK to create the Version1 version.
- n** In the Contents pane, notice that the connection reflects that you are now connected to the Version1 version, which is owned by publisher.



Next, you will use a different method to create another version.

- o From the Versioning tab, in the Versioning group, click Manage Versions.
- p Right-click the DEFAULT entry, which is the first row of the table, and choose New.
- q For Name, type **Version2** and press Enter.
- r For Version2, under Description, slowly click the entry two times, type **My second branch version**, and press Enter.
- s From the Versions tab, in the Manage Versions group, click Save to create the second version.

The connection still reflects that you are currently connected to the Version1 version.

- t Close the Versions: Cityboundary tab.

The Version Manager closes.

- u Close any open maps.
- v Save the project, and then exit ArcGIS Pro.
- w On the EBASE VM (blue desktop), close the web browser.

You successfully configured the Cityboundary data and service to leverage branch versioning. After sharing the layer with your ArcGIS Enterprise portal, you brought the Cityboundary web layer into a map in ArcGIS Pro and created two versions for editors to perform updates in isolated environments. Your organization is now ready to begin editing the branch-versioned data using the feature layer in ArcGIS Pro.

Lesson review

1. An analyst wants to publish a branch-versioned feature layer where the data is not copied and is maintained in a user-managed data store. Which type of user-managed data store does the analyst need to register?
 - a. Folder
 - b. File geodatabase
 - c. Cloud store
 - d. Enterprise geodatabase

2. A publisher has a file geodatabase that contains several rasters and an associated mosaic dataset. The publisher wants to share an imagery layer that references the mosaic dataset. What process should the publisher perform to accomplish this task?
 - a. Publish the imagery layer from ArcGIS Pro, then register the file geodatabase as a data store.
 - b. Register a separate raster data store with an ArcGIS Image Server site, then publish the imagery layer from ArcGIS Pro.
 - c. Copy the file geodatabase into an existing registered raster data store, then publish the imagery layer from ArcGIS Pro.
 - d. Register the folder containing the file geodatabase with an ArcGIS Image Server Site, then publish the imagery layer from ArcGIS Pro.

3. Registering a dataset as branch versioned requires that editor tracking be enabled.
 - a. True
 - b. False

4. Editing branch-versioned data requires access to a published feature service rather than a direct connection to the enterprise geodatabase.
 - a. True
 - b. False

Answers to lesson 5 questions

Exercise 5: Perform branch versioning workflows (page 5-11)

1. What is the database platform for this enterprise geodatabase?

PostgreSQL

2. Which additional options are automatically enabled?

The Archiving, Global IDs, and Editor Tracking options are required for branch versioning, so they are enabled when you choose to register a dataset as branch versioned.

In a previous lesson, you applied workload separation to isolate visualization and analysis by federating and configuring two ArcGIS Image Servers with ArcGIS Enterprise. This configuration decreased the number of system services running on each Image Server. You also learned that by taking advantage of web layers, you can also decrease the number of services running on ArcGIS Server. These best practices help limit the amount of system resources consumed by services.

In this lesson, you will build on previously learned best practices to understand how optimizing web services is important for overall system performance and the individual web service user experience. You will continue configuring additional service properties but this time to address performance, rather than enabling capabilities. You will explore service properties available through ArcGIS Server Manager and learn some service tuning troubleshooting techniques, including how to access and search ArcGIS Server logs.

Topics covered

Service instances

Working with shared instances

Resolving service performance issues

Service instances

Service instances are the mechanism that ArcGIS Server uses to make GIS resources available through URLs. When a client makes a request, such as to pan a map or navigate to an address, a service instance responds to that request using a Windows process called ArcSOC.exe.



Figure 6.1. In this example, a map service has a single service instance, executed by an ArcSOC process, to fulfill requests made to the service. The ArcSOC process can respond to one request at a time.

Because each ArcSOC process is allocated memory, there is a limit to the number of service instances that can be running on a machine at any given time. A sufficient number of ArcSOC processes are required to handle the traffic that your services receive. However, provisioning more instances than a service needs wastes computer memory and ultimately costs the organization time and money for not efficiently using server resources.

As an administrator, you have the ability to configure how ArcSOCs respond to requests from a client by configuring service properties. Service properties that impact the response of ArcSOC are grouped into four categories in ArcGIS Server Manager: pooling, parameters, processes and caching.

Service instances (continued)

Pooling

An application only uses an instance for the amount of time that it takes to complete one request, after which the instance is returned to the pool and is available to fulfill another request. If a service instance is available at the time that the request is made, the service instance immediately begins to fulfill the request. If an instance is not available at the time of the request, the request will be delayed or may even time out.



Figure 6.2. When a service instance is spun-up and started, it is ready to fulfill incoming requests from the client. When a request comes in, the ArcSOC responds to the request and sends the client back a result. When complete, the ArcSOC is released back into the pool and is ready to respond to additional requests.




Figure 6.3. If a client holds on to a service instance for longer than the maximum usage time, the service instance is automatically released, and the client loses its reference to the service instance.

Service instances (continued)




Figure 6.4. When all ArcSOC processes are in use, the request will queue until another service instance become available to handle the request.




Figure 6.5. When all ArcSOC processes are in use and a client's wait time exceeds the maximum wait time for a service, the request times out.

Parameters

Each type of service has associated parameters that can be configured to change how a service functions. Service parameters can include settings such as anti-aliasing (to balance performance and image quality) as well as properties that can assist with tuning services. Parameters will vary depending on the type of service. The following table includes two example parameter properties.

Service instances (continued)

| Parameters service property | Description |
|--|--|
| Maximum Number of Records Returned by Server | <ul style="list-style-type: none"> • Service property available for map services • The value controls the number of records returned from a service request • Increase if the service will be returning large numbers of records when fulfilling requests |
| Maximum image size per request | <ul style="list-style-type: none"> • Service property available for image services • The value controls the maximum number of pixel rows and columns returned by a request • Increase if the image service needs to display additional area/data |




Be cautious as these service properties will potentially increase the time needed for a service instance to fulfill the request and potentially overload the client.

Service instances (continued)

Processes

Processes contain service properties that specify how services will run as a process on the server. These properties are the same regardless of the type of service.

| Processes service property | Description |
|-----------------------------------|--|
| Service recycling intervals | <ul style="list-style-type: none"> • The cadence for the process of destroying and re-creating service instances on the ArcGIS Server machine to ensure that they are healthy and functioning • Modify the default value if you require a service to recycle at a different time of day <p> To reduce CPU usage spikes, consider staggering recycling intervals so that services are not all recycled at the same time(s).</p> |
| Health checks | Configure the frequency of health checks for service processes, such as checking and repairing data connections |

Service instances (continued)

Caching

Map and image services have the ability to be drawn dynamically or from a tile cache.

| Caching service property | Description |
|---------------------------|--|
| Dynamically from the data | The service is updated and drawn each time that a request is made to the service: <ul style="list-style-type: none"> You may elect to use dynamic services when data is updated frequently |
| Using tiles from a cache | The service distributes pre-rendered tiles at each request: <ul style="list-style-type: none"> You may elect to create a cache to reduce draw times and improve performance, especially if you maintain services to be used as basemaps When you enable caching, you are able to configure tiling schemes to determine the scales and level of detail for your tiles, as well as cache settings such as storage format and compression rates for tiles |

Service instance types

There are two types of service instances: dedicated and shared. Administrators should consider which of the following service instance types should be applied to each individual service to meet their organization's business needs.

Dedicated

Dedicated instances are dedicated or reserved for an individual service, meaning that the ArcSOCs will only respond to requests on behalf of one service. As an administrator or publisher of the service, you can specify the minimum and maximum number of service instances running at any given time. Services that receive constant traffic, are under a service level agreement, or receive compute-intensive requests may benefit from having their own dedicated server processes that are always available, even when the service is not being used.

Shared

Shared instances are pooled together and shared among multiple services, meaning that an ArcSOC could respond to a request from various services. Shared instances conserve memory usage by pooling several active server processes for use by multiple services. By doing so, it reduces the memory usage of services that are not actively handling requests. Because shared instances use systems resources more efficiently, services should be configured to use shared instances whenever possible.

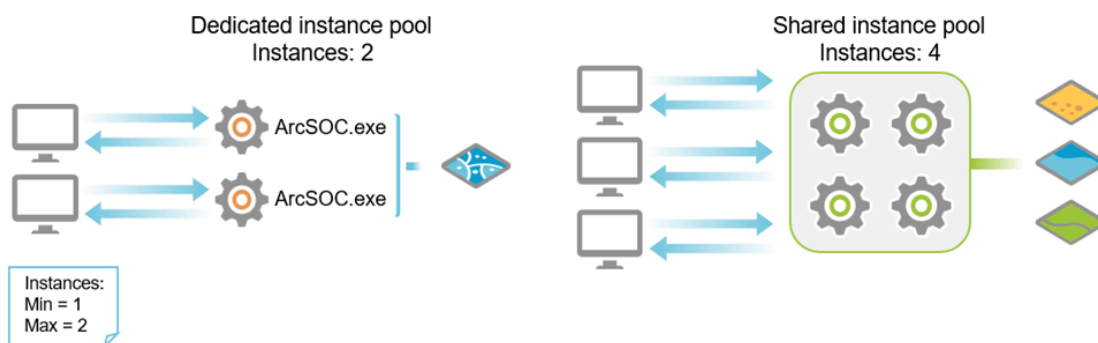


Figure 6.6. Dedicated instances are instances that are available to support requests to one service while shared instances are available to service requests to any service configured to use the shared instance pool.

Service instance types (continued)

Not all services support shared instances. When a service is published to your ArcGIS Server, it uses the default instance type. If the site was newly installed, the default instance type will be set to shared. If the site is upgraded, it will keep their current default instance setting. It is a best practice to have the default instance type set to use shared instances as it uses systems resources more efficiently. Dedicated and shared instances can be used together to optimize resource usage in your deployment.

Exploring service instances

When analyzing your services, you first need to gather baseline information about the service and its consumption of system resources. ArcGIS Server Manager and Windows Task Manager provide valuable information relating to your services and system resource usage. As you determine when to use shared and dedicated instances in your deployment, Task Manager can be used to examine the memory and CPU demands placed on your system. Windows Task Manager is a useful tool to view information about your ArcSOC processes and system resources. ArcGIS Server Manager is used to manage your services.

| Name | User name | CPU | Memory (p... | Command line |
|------------------------|-----------|-----|--------------|--|
| ArcGISCleanup.exe | adAdmin | 00 | 3,428 K | "C:\Program Files\ArcGIS\Pro\bin\ArcGISCleanup.exe" 11572 |
| ArcGISDataStore.exe | arcgisad | 00 | 194,096 K | "C:\Program Files\ArcGIS\DataStore\framework\etc\service\bin\ArcGISDataStore.exe" |
| ArcGISIndexingServ... | adAdmin | 00 | 42,436 K | "C:\Program Files\ArcGIS\Pro\bin\ArcGISIndexingServer.exe" /parentz:AOVerifier /ready={A9E56E93-383C-40E9-82B6-2C342BD5F058} |
| ArcGISPortal.exe | arcgisad | 00 | 304,524 K | "C:\Program Files\ArcGIS\Portal\framework\service\bin\ArcGISPortal.exe" |
| ArcGISPro.exe | adAdmin | 00 | 693,172 K | "C:\Program Files\ArcGIS\Pro\bin\ArcGISPro.exe" |
| ArcGISServer.exe | arcgisad | 00 | 221,548 K | "C:\Program Files\ArcGIS\Server\framework\etc\service\bin\ArcGISServer.exe" |
| ArcGISSharingServer... | adAdmin | 00 | 227,056 K | "C:\Program Files\ArcGIS\Pro\bin\ArcGISSharingServer.exe" ArcGISParent |
| ArcSOC.exe | arcgisad | 00 | 244,308 K | "C:\Program Files\ArcGIS\Server\framework\runtime\ArcGIS\bin\ArcSOC.exe" -XO: -CreateMinidumpOnCrash -Xms64M -XO: -UseParallelGC -Dservices\System.SpatialAnalysisTools.GPServerSync |
| ArcSOC.exe | arcgisad | 00 | 97,820 K | "C:\Program Files\ArcGIS\Server\framework\runtime\ArcGIS\bin\ArcSOC.exe" -XO: -CreateMinidumpOnCrash -Xms64M -XO: -UseParallelGC -Dservices\Utilities.GeocodingTools.GPServer -Djava |
| ArcSOC.exe | arcgisad | 00 | 106,936 K | "C:\Program Files\ArcGIS\Server\framework\runtime\ArcGIS\bin\ArcSOC.exe" -XO: -CreateMinidumpOnCrash -Xms64M -XO: -UseParallelGC -Dservices\System.DynamicMappingHost.MapServer |

Figure 6.7. In Windows Task Manager, enabling the Command line column provides the ability to identify which ArcSOC is tied to each service. Because they are not tied to a specific service, shared instances can be identified by the name DynamicMappingHost.

Consider when to use shared instances

As an administrator, you must determine which type of instance to use for your services to improve performance and optimize resource usage in your deployment. You will examine ArcGIS Enterprise Help documentation to answer questions regarding when to use shared instances.

Instructions

- a Open a web browser and go to the [ArcGIS Server: Configure Service Instance Settings](https://enterprise.arcgis.com/en/server/11.3/administer/windows/configure-service-instance-settings.htm) web page (https://enterprise.arcgis.com/en/server/11.3/administer/windows/configure-service-instance-settings.htm).
- b Review this page to answer the following questions, then close the web browser.



Read the sections titled *Shared and dedicated instances*, *When to use each instance type*, and *Configure the shared instance pool*.

1. Which types of services can use the shared instance pool?

2. Is a service that is infrequently used a good candidate for using shared instances?

Consider when to use shared instances (continued)

- 3. How do you determine the number of shared instances to configure for a given server host?**

Exercise 6

Work with shared instances

Imagine that you are an administrator, and you are responsible for using system resources effectively while also meeting user demands. One of your map services is used infrequently, and you believe that it may be a candidate for using the shared instance pool. In this exercise, you will inspect the implications of configuring this service to use shared instances rather than dedicated instances.

In this exercise, you will perform the following tasks:

- Use Windows Task Manager to inspect ArcSOC processes.
- Change the number of shared instances for a server site.
- Configure a service to use shared instances.

Step 1: Inspect ArcSOC processes using Task Manager


First, you will inspect the current processes running on your system to gain insight into resource usage.

- a Verify that you are working on the EBASE VM (blue desktop).
- b Right-click the Windows Taskbar and choose Task Manager.

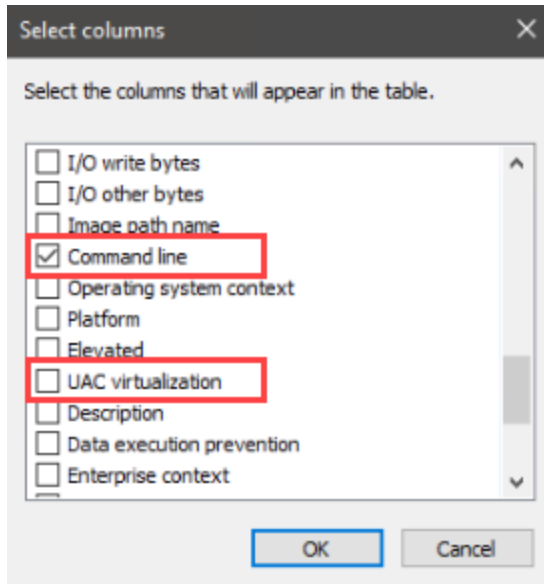


Task Manager is a useful application for viewing services and resource usage, but it should not be used to manage services, such as starting or stopping. ArcGIS Server Manager should be used for managing services.

- c In Task Manager, click the Details tab.
- d Right-click any column heading (such as Name) and choose Select Columns.



- e In the Select Columns window, uncheck the boxes for PID and Status.
- f Scroll down and locate the Command Line and UAC Virtualization check boxes.
- g Check the box for Command Line, and then uncheck the box for UAC Virtualization, as shown in the following graphic.



- h** Click OK.
- i** Adjust the column widths so that you can see the contents of the Command Line column.
- j** Click the Command Line column heading to sort the processes.
- k** Scroll down and locate the processes named ArcSOC.exe.

| Name | User n... | C... | Memory... | Command line |
|-------------------------------------|-----------|------|-----------|--|
| <input type="checkbox"/> ArcSOC.exe | arcgis... | 00 | 115,652 K | "C:\Program Files\ArcGIS\Server\framework\runtime\ArcGIS\bin\ArcSOC.exe" -XX:-Cr |
| <input type="checkbox"/> ArcSOC.exe | arcgis... | 00 | 114,704 K | "C:\Program Files\ArcGIS\Server\framework\runtime\ArcGIS\bin\ArcSOC.exe" -XX:-Cr |
| <input type="checkbox"/> ArcSOC.exe | arcgis... | 00 | 115,992 K | "C:\Program Files\ArcGIS\Server\framework\runtime\ArcGIS\bin\ArcSOC.exe" -XX:-Cr |
| <input type="checkbox"/> ArcSOC.exe | arcgis... | 00 | 116,140 K | "C:\Program Files\ArcGIS\Server\framework\runtime\ArcGIS\bin\ArcSOC.exe" -XX:-Cr |
| <input type="checkbox"/> ArcSOC.exe | arcgis... | 00 | 114,452 K | "C:\Program Files\ArcGIS\Server\framework\runtime\ArcGIS\bin\ArcSOC.exe" -XX:-Cr |
| <input type="checkbox"/> ArcSOC.exe | arcgis... | 00 | 115,304 K | "C:\Program Files\ArcGIS\Server\framework\runtime\ArcGIS\bin\ArcSOC.exe" -XX:-Cr |

- l** For the processes named ArcSOC.exe, inspect the Command Line column and identify the entries that include -Dservice=System.DynamicMappingHost.MapServer, as shown in the following graphic.

```

+UseParallelGC -Dservice=System.DynamicMappingHost.MapServer --add-
+UseParallelGC -Dservice=System.DynamicMappingHost.MapServer --add-
+UseParallelGC -Dservice=System.DynamicMappingHost.MapServer --add-
+UseParallelGC -Dservice=System.DynamicMappingHost.MapServer --add-
    
```

Hint: You may have to scroll to the right to find this portion of the Command Line entry.

These ArcSOC processes correspond to the number of shared instances set for your site.

1. Based on what you see in Task Manager, how many shared instances are configured for the GIS Server site configured on the EBASE machine?
-

- m** Inspect the Command Line column and identify the entries that include -Dservice=Manhattan_Streets.MapServer, as shown in the following graphic.

```
GC -Dservice=Manhattan_Streets.MapServer --add-  
GC -Dservice=Manhattan_Streets.MapServer --add-  
GC -Dservice=Manhattan_Streets.MapServer --add-  
GC -Dservice=Manhattan_Streets.MapServer --add-  
GC -Dservice=Manhattan_Streets.MapServer --add-  
GC -Dservice=Manhattan_Streets.MapServer --add-
```

2. Based on what you see in Task Manager, how many dedicated instances are configured for the Manhattan_Streets service?
-

3. How much memory is being consumed in total by all the Manhattan_Streets ArcSOC processes?
-

- n** Leave Task Manager open.

In this step, you added the Command Line column in Task Manager to identify both dedicated and shared service instances. You observed that each ArcSOC consumes memory.

Step 2: Change shared instance settings for a server site


Next, you will change the number of shared instances configured for your site based on best practices.

- a On the EBASE VM (blue desktop), open a web browser, and then click the EBASE-Server Manager bookmark (<https://ebase.ad.local/server/manager>).
- b Sign in as **portaladmin** with the **Esri.4.GIS** password, if necessary.
- c Click the Site tab, and then click the Settings tab.

4. What is the default instance type for this site?

You will keep shared instances as the default service type for services published to this site.

5. After checking with your IT department, you learn that the EBASE machine is installed on a virtual machine with six cores. You will be using an equal mix of dedicated and shared instances. How many shared instances should you configure?

- d Next to Shared Instance Settings, click the Edit Site Mode button .
- e In the Shared Instance Settings window, for Number Of Shared Instances Per Machine, type **3**, and then click Apply.

Shared instance settings are specific to each server site.

- f Restore the Task Manager window.
- g In Task Manager, verify that there are now three processes used for shared instances.

Hint: Look for entries in the Command Line column that include
 -Dservice=System.DynamicMappingHost.MapServer.


In this step, you edited the number of shared service instances running on each ArcGIS Server machine in the site.

Step 3: Configure a service to use shared instances

You will now configure the Manhattan_Streets service to use shared instances. Because this map service is infrequently used, you believe that switching to shared instances will help optimize resource usage. After making the change, you will observe the impacts to system resource usage to confirm that the service is consuming less memory.

- a Restore the web browser tab with ArcGIS Server Manager.
- b Click the Services tab, and then click the Manage Services tab, if necessary.

- c Click the link for the Manhattan_Streets service.
- d Click the Pooling tab.
- e Under Instance Type, for Make This Service Available Using, select Shared Instance Pool.



- f Click Save And Restart to restart the service.
- g Next to the Manhattan_Streets service link, click the thumbnail, as shown in the following graphic.



A new web browser tab opens to a map that visualizes the service using the ArcGIS API for JavaScript.

- h Sign in as **portaladmin** with the **Esri.4.GIS** password, if necessary.
- i Resize the web browser window so that it covers approximately half of the screen.
- j In the other half of the screen, position the Task Manager window so that the CPU and Memory columns for the three shared instance processes are visible.

Hint: You can click one of the rows that correspond to a shared instance to highlight it, which can aid in keeping track of the row.

- k While observing the shared instance processes in Task Manager, zoom in several levels and pan the map in the JavaScript viewer.

6. What do you observe in Task Manager?

7. Are there any processes that correspond to dedicated instances for the Manhattan_Streets service?

- l** Close Task Manager.
- m** Close the ArcGIS API for JavaScript web browser tab.

In this step, you configured the Manhattan_Streets map service to leverage shared instances rather than dedicated.

Step 4: Explore deletion behavior of a non-hosted feature layer

You will now explore deletion behavior of user-managed data as it relates to layers and services. First you will enable Feature Access for the Manhattan_Streets map service.

- a** In ArcGIS Server Manager, click the link for the Manhattan_Streets map service.
- b** Click the Capabilities tab.
- c** Under Select And Configure Capabilities, check the box for Feature Access.

Select and configure capabilities

| | |
|--|--|
| <input checked="" type="checkbox"/> Mapping (always enabled) | <input type="checkbox"/> Validation |
| <input type="checkbox"/> WCS | <input type="checkbox"/> WFS |
| <input type="checkbox"/> Network Analysis | <input type="checkbox"/> Topographic Production |
| <input type="checkbox"/> Version Management | <input type="checkbox"/> Network Diagrams |
| <input type="checkbox"/> Trace Network | <input type="checkbox"/> KML |
| <input type="checkbox"/> Linear Referencing | <input type="checkbox"/> WMS |
| <input type="checkbox"/> OGC Features | <input type="checkbox"/> Utility Network |
| <input type="checkbox"/> Parcel Fabric | <input checked="" type="checkbox"/> Feature Access |

d Click Save And Restart.

By enabling Feature Access, a feature service is now exposed for Manhattan_Streets.

e Open a new web browser tab, and then click the ArcGIS Enterprise bookmark (<https://ebase.ad.local/portal/home>).

f Sign in using the portaladmin account (password: **Esri.4.GIS**), if necessary.

g Click Content, and then click the My Organization tab.

h In the Search My Organization field, type **title: Manhattan_Streets** and press Enter.

There are now two layers exposed in the portal for the Manhattan_Streets service: a feature layer and a map image layer.

i Check the box for the Manhattan_Streets feature layer.


| Title | | Modified |
|---|-----------------|--------------|
| <input type="checkbox"/> Manhattan_Streets | Map Image Layer | May 25, 2022 |
| <input checked="" type="checkbox"/> Manhattan_Streets | Feature Layer | Jun 13, 2024 |

j Click Delete, and then click Delete to confirm the deletion.

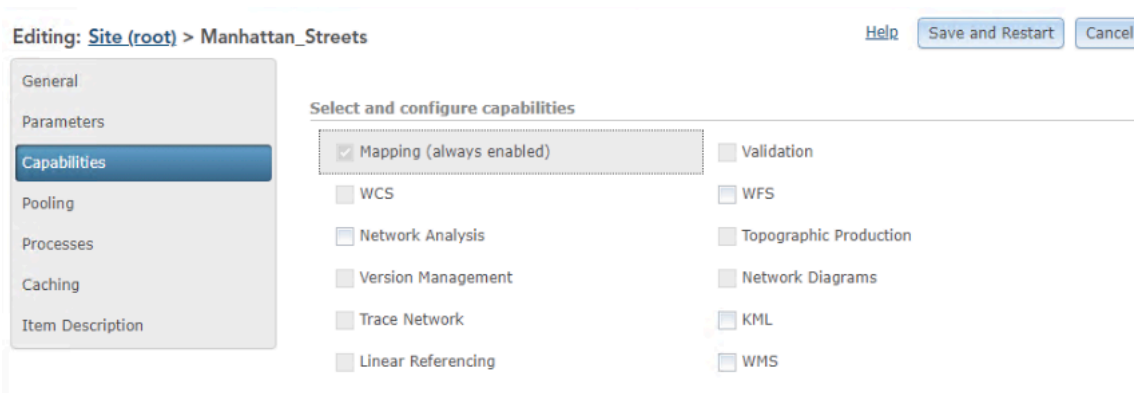
Even though you deleted the feature layer from the Enterprise portal, the map image layer created during publishing still exists.



If you attempt to delete the map image layer before the feature layer, you will instead receive an error message that says that you cannot delete it due to dependent items, as shown in the following graphic. If you need to delete both portal items, you must first delete the feature layer, and then you can delete the map image layer.



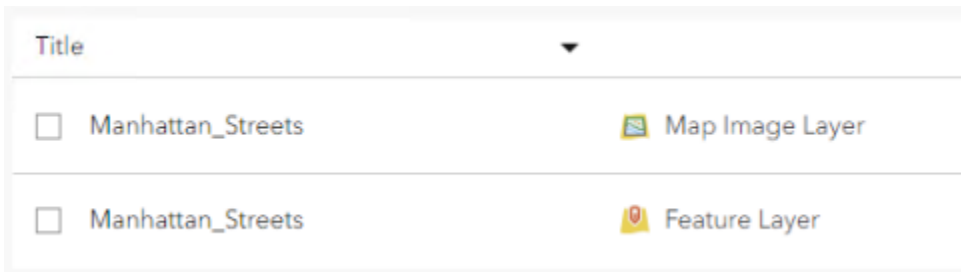
- k** Restore the web browser tab that is signed in to ArcGIS Server Manager (<https://ebase.ad.local/server/manager>).
- l** From the main Services tab, click the Manage Services tab, and then under Folders, click Site (Root).
- m** Notice that the Manhattan_Streets map service still exists in ArcGIS Server Manager.
- n** Click the Manhattan_Streets map service, and then click the Capabilities tab.



As a result of deleting the feature layer item from Enterprise portal, the Manhattan_Streets map

service no longer has feature access capabilities. While you can still utilize the map service for mapping capabilities, you can also re-enable feature access capabilities to the map service and restore the Manhattan_Streets feature layer in the Enterprise portal.

- o Under Select And Configure Capabilities, check the box for Feature Access, and then click Save And Restart.
- p Restore the web browser tab that is signed in to the Enterprise portal.
- q Refresh the browser.



After you re-enable feature access capabilities and restart the map service, you will automatically restore the feature layer that you had previously deleted from the Enterprise portal. As an additional benefit to working with feature layers and map services published from data within a referenced data store, any edits that are made to the data will be retained in the original data source and will not be lost during this deletion and restoration process.

- r Close the web browser.

In this exercise, you worked with dedicated and shared instances and inspected the resource usage of their corresponding ArcSOC processes. You also explored deletion behavior specific to services that use user-managed data.

Reviewing ArcGIS Server logs and statistics reports

ArcGIS Server logs can be useful in several ways:

- Getting statistics on service performance
- Troubleshooting problems with an ArcGIS Server
- Documenting user workflows

Reviewing ArcGIS Server logs and statistics reports (continued)



Figure 6.8. Statistic reports can be configured to help detect when and how a service should be tuned for a better user experience and for a responsible use of system resources.

Administrators can specify the number of days before logs are deleted (the default is 90 days), as well as set the desired logging level for ArcGIS Server. For production servers, the logging level should always be set to Severe or Warning after the investigation of an issue is finished.

Reviewing ArcGIS Server logs and statistics reports (continued)



ArcGIS Enterprise Help: *About specifying server log settings*

ArcGIS Enterprise Help: *Work with server logs*

ArcGIS Enterprise Help: *Log codes overview*

Troubleshooting common service performance issues

Leveraging tools such as ArcGIS Server logs or ArcGIS Monitor can help determine when services need tuning. Service properties can be modified to improve the performance of your services. Based on user load, service functionality, and other factors, you may need to adjust the properties away from the default values.

For each of the following scenarios, use the information found in previous pages of the workbook to determine which property should be modified and describe how you would modify it.



ArcGIS Server Help: *Tune services using best practices*

Scenario 1: Service with initial use delays

You have a popular service that experiences relatively constant demand throughout the day. The initial users of the service in the early morning complain that the service is slow to respond. You have never received any complaints about this service performance throughout the rest of the day.

- 1. Which service property should you modify and how should it be set?**
-
-

Scenario 2: Numerous services with limited system resources

As the GIS coordinator, you have grown the adoption of ArcGIS Enterprise across the organization. Users have been publishing many map services from ArcGIS Pro that use dedicated instances, increasing the load on already limited system resources. An audit shows that many of these map services are used only sporadically throughout the month.

Troubleshooting common service performance issues (continued)

2. Which service property should you modify and how should it be set?

Scenario 3: Incomplete result sets from a geoprocessing service

Users are reporting that sometimes after successfully running a geoprocessing service, a query on the output is not returning all the results. They have indicated that everything seems to work fine when they are expecting a small result set, but the issue occurs for larger result sets.

3. Which service property should you modify and how should it be set?

Scenario 4: Service experiencing increased demand

You have noticed in a server log that a service that was recently mentioned in the company newsletter is now experiencing regular wait-time timeouts.

4. Which service property should you modify and how should it be set?

Lesson review

1. Which type of service may be a candidate for using the shared instance pool rather than dedicated instances?
 - a. A service for which the minimum dedicated instances is set to zero
 - b. A service receiving nearly constant, high-volume demand
 - c. A service operating under a service-level agreement (SLA)
 - d. An infrequently used geoprocessing service
2. Users have reported waiting a long time to access a particular service. Which service property must be increased to resolve the issue?
 - a. Minimum number of instances per machine
 - b. Maximum number of instances per machine
 - c. Maximum time that a client can use a service
 - d. Maximum time that a client will wait to get a service
3. Which server statistic can be tracked using reports in ArcGIS Server Manager?
 - a. Originating IP addresses of requests
 - b. Maximum running instances
 - c. Fastest response time
 - d. Usernames associated with requests

Answers to lesson 6 questions

Consider when to use shared instances (page 6-11)

1. Which types of services can use the shared instance pool?

Map services published from ArcGIS Pro can be configured for the shared instance pool. Capabilities other than feature access, WFS, WMS, and KML must be turned off before publishing.

2. Is a service that is infrequently used a good candidate for using shared instances?

Yes; services that are infrequently used tend to be good candidates for shared instances (for example, services for which you have already set the minimum dedicated instances to zero).

3. How do you determine the number of shared instances to configure for a given server host?

If most or all of your site's services use the shared instance pool, set the number of shared instances to twice the number of physical CPU cores on the machine. If you are using hyperthreading, vCPUs, or logical cores (such as if your site is in a virtual or cloud environment), consider each logical core equal to one-half of a physical core.

Exercise 6: Work with shared instances (page 6-13)

1. Based on what you see in Task Manager, how many shared instances are configured for the GIS Server site configured on the EBASE machine?

4

2. Based on what you see in Task Manager, how many dedicated instances are configured for the Manhattan_Streets service?

6

Answers to lesson 6 questions (continued)

3. How much memory is being consumed in total by all the Manhattan_Streets ArcSOC processes?

Several hundred MB

4. What is the default instance type for this site?

Shared

5. After checking with your IT department, you learn that the EBASE machine is installed on a virtual machine with six cores. You will be using an equal mix of dedicated and shared instances. How many shared instances should you configure?

The number of shared instances should be equal to the number of virtual cores on the machine, divided in half: three.

6. What do you observe in Task Manager?

The CPU and memory usage increases slightly for one of the shared instance processes.

7. Are there any processes that correspond to dedicated instances for the Manhattan_Streets service?

No; this map service now fully uses the shared instance pool.

Troubleshooting common service performance issues (page 6-26)

Scenario 1: Service with initial use delays

1. Which service property should you modify and how should it be set?

Increase the minimum instances value to ensure that more instances will be ready first thing in the morning.

Scenario 2: Numerous services with limited system resources

2. Which service property should you modify and how should it be set?

Switch the service instance type to shared to free up system resources by limiting the number of ArcSOCs running on the ArcGIS Server machine.

Answers to lesson 6 questions (continued)

Scenario 3: Incomplete result sets from a geoprocessing service

3. Which service property should you modify and how should it be set?

Increase the maximum returned records value to a point where all requests will get complete result sets.

Scenario 4: Service experiencing increased demand

4. Which service property should you modify and how should it be set?

Increase the maximum instances to accommodate the increased demand.

You have been introduced to the two types of data in ArcGIS Enterprise: user-managed and ArcGIS-managed. At the core of ArcGIS-managed data is ArcGIS Data Store, which configures and manages data storage for the hosting server. In doing so, ArcGIS Data Store allows your ArcGIS Enterprise portal users to publish large numbers of hosted feature layers, hosted scenes (3D), archive high-volume real-time observational data, store knowledge graphs, and cache queries for hosted feature layers.

ArcGIS Data Store is administered primarily by ArcGIS Enterprise, but administrators have access to various utilities to perform key administrative functions, such as managing backups. Backups can be created for most ArcGIS Data Store types. As an administrator, you have various options for specifying the backup location, frequency, and retention schedule.

Topics covered

- ArcGIS Data Store functions and types

- Hosted and non-hosted services

- Managing ArcGIS Data Store using command utilities

- ArcGIS Data Store backups

ArcGIS Data Store

ArcGIS Data Store is a software application that creates and manages data storage for ArcGIS Enterprise. There are numerous types of data stores that can be configured by ArcGIS Data Store, all of which are configured with the hosting server that is a part of the ArcGIS Enterprise base deployment. Each type provides key functionality to your users.

ArcGIS Data Store types

ArcGIS Data Store can be set up in the following ways to provide functionality to your system:

- The **relational data store** stores feature data for your hosted feature layers.
- The **tile cache data store** stores caches for hosted scene layers.
- The **spatiotemporal big data store** stores high-volume observational data that often includes both space and time components.
- The **graph store** stores relationships between entities in a knowledge graph.
- The **object store** caches query responses for hosted feature layers for improved rendering performance.




Figure 7.1. ArcGIS Data Store can create five different types of data stores, all of which are always configured with the hosting server.

Both the relational and tile cache data stores are required as part of a base ArcGIS Enterprise deployment, and they may be configured together on one machine or on separate machines. Spatiotemporal big data stores, graph stores, and object stores are not required for a base deployment and must be installed on a dedicated host.

ArcGIS Data Store (continued)



ArcGIS Data Stores that are configured on the same machine compete for memory and other resources, which can negatively affect performance. Esri recommends that you install data store types on separate machines to avoid resource contention between data stores, which is especially important in the case of spatiotemporal big data stores, graph stores, and object stores. Do not install these data stores on the same machine as other data stores or ArcGIS Enterprise components.

Hosting server

When configured, the relational data store allows the ArcGIS GIS Server site to be designated in the Enterprise portal as the *hosting server*. Similar to user-managed data stores, data stores managed by ArcGIS Data Store are always configured with an ArcGIS GIS Server site. Unlike user-managed data stores, however, data stores managed by ArcGIS Data Store will only be configured with the hosting server. All services with underlying data stored in a data store managed by ArcGIS Data Store will be hosted on the hosting server.

ArcGIS GIS Server
(hosting server)



ArcGIS Data Store



Figure 7.2. The ArcGIS Data Store is one of the required components that make up the ArcGIS Enterprise base deployment, and data stores managed by ArcGIS Data Store are always configured with the hosting server.

Types of ArcGIS-managed data

Within the two broad categories of data sources in ArcGIS (user-managed and ArcGIS-managed), there are several data types that you can use to meet your business needs. Having already reviewed user-managed data, the following list provides examples of six different types of ArcGIS-managed data.

- **Hosted feature layers** result from publishing specific types of data or by feature analysis tasks.
- **Hosted imagery layers** result from rasters uploaded as an image collection or result from data generated by raster analysis tasks.
- **High-volume observation data** result from potentially massive datasets of streamed real-time data generated by GeoEvent Server or locations recorded from apps such as ArcGIS Tracker.
- **Hosted scene layers** result from publishing 3D object, building, integrated mesh, point, and point cloud layers.
- **Cached query responses** result from caching query responses from hosted feature layers.
- **Knowledge graphs** result from storing relationships created from graph analysis.

In the following graphic, identify which example of ArcGIS-managed data is stored within each ArcGIS Data Store type.

- A = High-volume observation data
- B = Hosted feature layers
- C = Hosted imagery layers
- D = Hosted scene layers (3D)
- E = Cached query responses
- F = Knowledge graphs

Types of ArcGIS-managed data (continued)

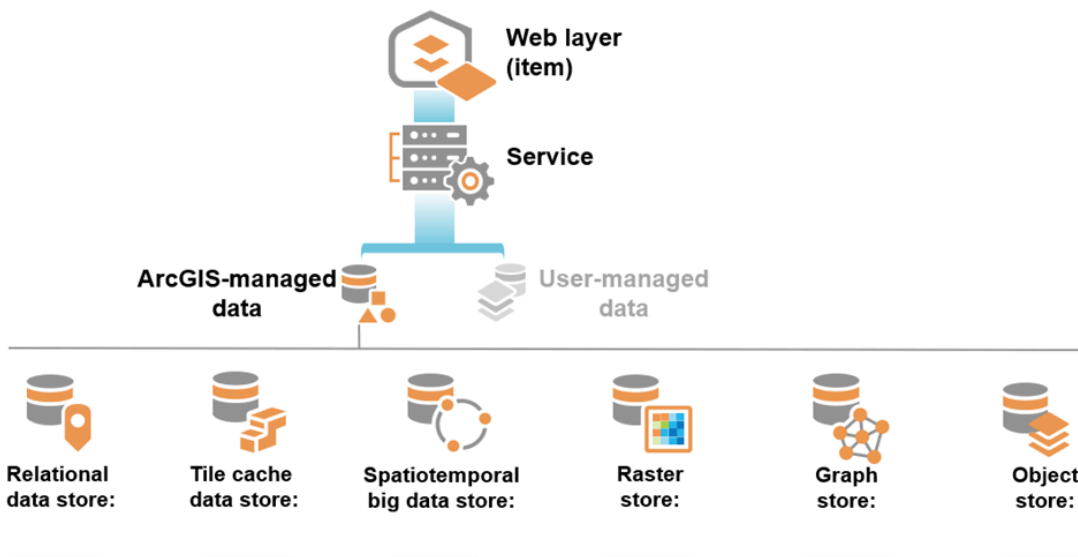


Figure 7.3. Match the underlying data type associated with each type of ArcGIS-managed data store.

1. Relational data store

2. Tile cache data store

3. Spatiotemporal big data store

4. Raster store

5. Graph store

6. Object store

Relationship between data storage options and service type

When publishing from ArcGIS Pro, you determine the location of the data used by your web layers and corresponding services. For most data sources, you can either register your data store—in which case the web layers access the data in place (user-managed)—or have ArcGIS Enterprise copy the data to a location managed by ArcGIS Enterprise (ArcGIS-managed).

Hosted and non-hosted services

There is often a misconception that if a web layer's underlying data is ArcGIS-managed, then the layer and associated service is always considered hosted. Additionally, it is often thought that user-managed data means that the layer and service are non-hosted; however, that is not always the case. The distinction between hosted and non-hosted services is important to understand because it will determine whether the service will leverage service instances, as well as determining other services properties that can be configured.

- Hosted services are managed by ArcGIS Enterprise and do not leverage service instances.
- Non-hosted services require tuning and management of service instances by an administrator.

The following table lists types of layers and whether they can be hosted or non-hosted and whether they use ArcGIS-managed or user-managed data.

Relationship between data storage options and service type (continued)

| | ArcGIS-managed data | User-managed data |
|------------|--|---|
| Hosted | <ul style="list-style-type: none"> • Feature layer • Map image layer • Tile layer • Vector tile layer • Imagery layer • Scene layer • Knowledge graph | <ul style="list-style-type: none"> • Vector tile layer |
| Non-hosted | <ul style="list-style-type: none"> • Map image layer • Imagery layer | <ul style="list-style-type: none"> • Feature layer • Map image layer • Imagery layer |

ArcGIS Data Store backup options

Although layers and services that use ArcGIS-managed data are stored and managed by ArcGIS Enterprise, administrators still need to consider a backup strategy for this data storage option. Backups allow you to recover your data if a disaster occurs, such as when your server fails or floodwaters destroy your server. There are two ways to create backups of the data contained in ArcGIS Data Store.

1. ArcGIS Data Store command line utilities support the configuration and running of both full and incremental backups of data contained in most ArcGIS Data Store types. There are ArcGIS Data Store commands to identify backup data locations, frequency of automatic backups, and the length of time that backups are to be retained. It is also possible to execute a command to perform a manual backup. If a backup is done manually, the backup will be retained indefinitely.
2. The ArcGIS Enterprise deployment backup utility, *webgisdr*, can also back up the data in the relational and tile cache data stores. However, the *webgisdr* utility is more comprehensive. Use the *webgisdr* utility to export backup files of the following components of your ArcGIS Enterprise deployment:
 - Your portal items and settings
 - GIS services and settings
 - The relational data store and tile cache data store



If your backup files are located on the server that is destroyed in a disaster, such as a flood, then you cannot recover your data; therefore, you should store your backup files on a different server than your data store.

ArcGIS Data Store backup defaults

When considering a backup strategy for ArcGIS Data Store, it is important to note the various backup properties. All ArcGIS Data Stores can be backed up to the following locations:

- File share
- Amazon S3 bucket
- Azure blob

However, the default backup properties vary depending on the type of ArcGIS Data Store.

The following table describes the default backup properties for each ArcGIS Data Store type.

| ArcGIS Data Store | Default backup location | Default backup schedule |
|---------------------------|--|--------------------------------------|
| Relational data store | On the machine where the relational data store is installed | Every four days, retained seven days |
| Tile cache data store | <ul style="list-style-type: none"> • On the machine where the tile cache data store is installed • No default when created in cluster mode | None |
| Spatiotemporal data store | None | None |
| Graph store | None | None |
| Object store | None | None |

Relational data store

When you create a relational data store, a backup location is automatically configured on the same machine as the data store. By default, ArcGIS Data Store creates a full backup of the relational data store every four days and saves them for seven days in **C:\arcgisdatastore\backup\relational**. If the primary data store machine fails and you have not moved the backup location

ArcGIS Data Store backup defaults (continued)

from the default path, you will not be able to restore your hosted feature layer data.

Tile cache data store

The tile cache data stores created in primary-standby mode have a default backup location while tile cache data stores created in cluster mode do not. When you create a tile cache data store in cluster mode, you must register at least one external backup location before you can create tile cache data store backup files. Like relational data stores, tile cache data stores can be backed up to a shared file directory or a cloud storage container.

Spatiotemporal big data store

Spatiotemporal big data stores do not have a default backup location. You must register at least one backup location before you can create spatiotemporal big data store backup files.

Although it is possible to configure automatic backups for the spatiotemporal big data store, the data store itself is designed to support both parallel data writing and redundancy. As additional spatiotemporal data store nodes are added to a portal configuration, the data is distributed across the nodes in such a way that the failure of any single node will not result in a loss of data.

Graph store

Similar to the spatiotemporal big data store, the graph store does not have a default backup location or default backup schedule. However, to back up the graph store, the same utility (`backupdatastore`) can be used.



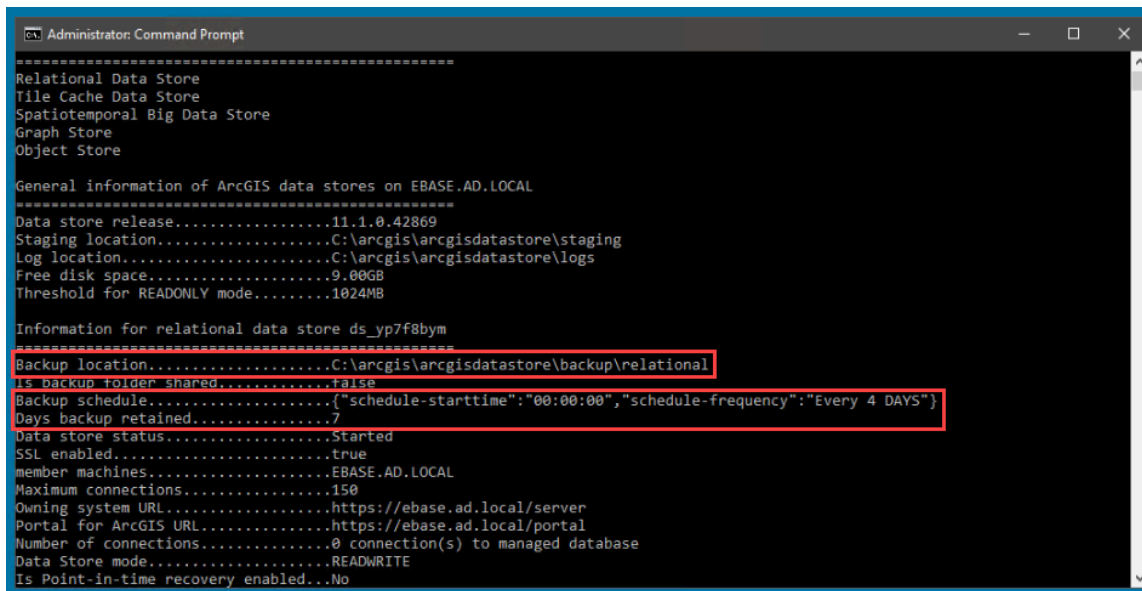
Be aware that there is no automatic cleanup of tile cache, spatiotemporal big data store, or graph store backup files. If you enable automatic backups for these data stores, you will likely need to clean up the backup location as time passes.

Object store

At this time, backups are not supported for the object store.

Managing ArcGIS Data Store using command utilities

ArcGIS Data Store provides many utilities to perform management tasks. The utilities are executed using the Command Prompt on the ArcGIS Data Store machine.



```

Administrator: Command Prompt
=====
Relational Data Store
Tile Cache Data Store
Spatiotemporal Big Data Store
Graph Store
Object Store

General information of ArcGIS data stores on EBASE.AD.LOCAL
=====
Data store release.....11.1.0.42869
Staging location.....C:\arcgis\arcgisdatastore\staging
Log location.....C:\arcgis\arcgisdatastore\logs
Free disk space.....9.00GB
Threshold for READONLY mode.....1024MB

Information for relational data store ds_yp7f8bym
=====
Backup location.....C:\arcgis\arcgisdatastore\backup\relational
Is backup folder shared.....false
Backup schedule.....{"schedule-starttime":"00:00:00","schedule-frequency":"Every 4 DAYS"}
Days backup retained.....7
Data store status.....Started
SSL enabled.....true
member machines.....EBASE.AD.LOCAL
Maximum connections.....150
Owning system URL.....https://ebase.ad.local/server
Portal for ArcGIS URL.....https://ebase.ad.local/portal
Number of connections.....0 connection(s) to managed database
Data Store mode.....READWRITE
Is Point-in-time recovery enabled...No
  
```

Figure 7.4. This example shows the result of executing the `describedatastore` command utility.

The following table identifies several common administrative task categories and some associated command utilities.

Managing ArcGIS Data Store using command utilities (continued)

| Category | Example command utilities |
|--|---|
| Obtain information about ArcGIS Data Store | <ul style="list-style-type: none"> • <i>describedatastore</i> <ul style="list-style-type: none"> • Obtain a variety of information about all data stores on a machine. • <i>listadminusers</i> <ul style="list-style-type: none"> • Obtain a list of the usernames and passwords for administrative accounts. |
| Manage certificates | <ul style="list-style-type: none"> • <i>updatesslcertificate</i> <ul style="list-style-type: none"> • Replace the self-signed certificate for any type of data store with a CA-signed or domain certificate. |
| Manage a highly available system | <ul style="list-style-type: none"> • <i>remove machine</i> <ul style="list-style-type: none"> • Remove a standby machine from a relational data store or a machine from a tile cache data store or spatiotemporal data store. |

Table continued on next page

Managing ArcGIS Data Store using command utilities (continued)

| Category | Example command utilities |
|----------------|--|
| Manage backups | <ul style="list-style-type: none">• <i>backupdatastore</i><ul style="list-style-type: none">• Perform a manual backup of a data store.• <i>listbackups</i><ul style="list-style-type: none">• See the names of backup files, the location to which they are written, and additional details about completed or ongoing backups.• <i>updatebackupschedule</i><ul style="list-style-type: none">• Change how often a full backup is created.• <i>configurebackuplocation</i><ul style="list-style-type: none">• Specify the location where ArcGIS Data Store writes backup files; this command also allows you to remove a backup location. |



Some command utilities may only apply to certain ArcGIS Data Store types. The ArcGIS Enterprise Help documentation that is listed below provides additional details.



ArcGIS Enterprise Help: ArcGIS Data Store command utility reference

Exercise 7

Manage ArcGIS Data Store backups

Imagine that you are the administrator who is responsible for backing up the data stored in the relational ArcGIS Data Store. You want to use the best practice of configuring a backup location on a separate machine using a shared network folder. To do this action, you plan on creating a folder on the IMAGE machine, which you have determined has sufficient resources for storing the backups. Then, you will use ArcGIS Data Store's command utilities to manage the backup settings for the relational data store installed on the EBASE machine.

In this exercise, you will perform the following tasks:

- Create a new folder to contain the backups.
- Use command line utilities to configure relational ArcGIS Data Store backup options.
- Perform a backup of the relational data store.

Step 1: Create a shared network folder

Your first task is to create a new shared network folder to which ArcGIS Data Store will write backups. You will create this folder on the IMAGE machine.

- a Connect to the IMAGE VM (orange desktop) using the Remote Desktop Connection Manager.

You will create the new folder on the C: drive of the IMAGE machine.

- b From the Windows Taskbar, open File Explorer.
- c In File Explorer, under This PC, right-click Local Disk (C:), point to New, and choose Folder.
- d Rename the folder **RelationalDataStoreBackups**.

You must now configure the sharing properties to allow ArcGIS Data Store to write to this folder.

- e Right-click the RelationalDataStoreBackups folder and choose Properties.
- f In the RelationalDataStoreBackups Properties dialog box, click the Sharing tab.

1. What is displayed for the Network Path of this folder?


- g Click Share.

You must share this folder with the `arcgis$` group-managed service account, which is the account that runs the ArcGIS Data Store Window's service and needs read/write permissions to the new folder.

- h In the Network Access dialog box, under Choose People On Your Network To Share With, type **arcgis\$** and click Add.
- i Right-click `AD\arcgis$` and choose Read/Write.

| Name | Permission Level |
|----------------------------|------------------|
| AD\arcgis\$ | Read/Write ▼ |
| adAdmin (adAdmin@ad.local) | Owner |

- j Click Share.
- k Notice the UNC file path that is displayed.




- l Click Done.
- m Open the RelationalDataStoreBackups Properties dialog box and click the Sharing tab, if necessary.

2. In the RelationalDataStoreBackups Properties dialog box, what is displayed for the Network Path now?

- n Highlight the network path (\\IMAGE\RelationalDataStoreBackups), right-click the highlighted text, and choose Copy.
- o Connect to the EBASE VM (blue desktop).

As a test, you will attempt to access this shared folder from the EBASE machine.

- p In the EBASE VM, open File Explorer, and then in the address bar, paste the network path to the shared folder.



- q Press Enter.

Because you shared this folder, it is now accessible from another machine in the network.

Step 2: Run a command utility

Now you can begin configuring the ArcGIS Data Store settings. First, you will use the command utilities to get some basic descriptive information about ArcGIS Data Store.

- a On the EBASE VM (blue desktop), in File Explorer, browse to **C:\Program Files\ArcGIS\DataStore\tools**.

This folder contains the ArcGIS Data Store command utilities. You will reference this folder location when preparing to run utilities using the Command Prompt.

- b** Click the address bar to highlight the folder path (**C:\Program Files\ArcGIS\DataStore\tools**), right-click the highlighted text, and choose Copy.

You will now open the Command Prompt with administrator privileges and run a command utility.

- c** From the Windows Taskbar, right-click the Windows Start icon and choose Command Prompt (Admin).
- d** In the Command Prompt, type `cd`, press the spacebar, and then paste the path that you copied.

```
C:\WINDOWS\system32>cd C:\Program Files\ArcGIS\DataStore\tools_
```

This command will change your current working directory to the folder where the ArcGIS Data Store command utilities are located.

- e** Press Enter to run the command.

With your working directory set, you can execute command utilities by simply typing their name.

First, you will obtain some descriptive information about ArcGIS Data Store using the `describedatastore` utility.

- f** In the Command Prompt, type `describedatastore.bat`.

Hint: In the Command Prompt, you can partially type the name of a command utility, such as `desc`, and then press the Tab key to have Command Prompt complete the name based on the file names within your working directory. Continuing to press Tab will cycle through any other available matches.

- g** Press Enter to execute the `describedatastore` utility.

The output contains a variety of key information about ArcGIS Data Store, including settings related to backups.

3. Where is the current backup location for the relational data store?

The backup location will be stored on the same machine as ArcGIS Data Store, which makes ArcGIS Data Store vulnerable in the event that this machine becomes unavailable.

4. Is point-in-time recovery enabled?

You must enable point-in-time recovery if you will use the `webgisdr` utility to create incremental backups of your entire ArcGIS Enterprise deployment. If you enable point-in-time recovery, incremental backups are created either when the log files are full or every five minutes, whichever comes first. The database controls incremental backup creation; you cannot control the frequency with which incremental backups are created.

Step 3: Configure relational data store backup properties

Next, you will change the backup location. You will first place a relational data store in read-only mode, which is a best practice, while you perform maintenance on ArcGIS Data Store.

- a In the Command Prompt, type `changedatastoremode readonly --store relational` and press Enter.
- b When you see the prompt asking whether you want to continue, type `yes` and press Enter.

With ArcGIS Data Store in read-only mode, you can safely change settings while ensuring data integrity. For example, you do not want users to be in the process of publishing content while administrative work is happening.

- c Type `configurebackuplocation.bat --help` and press Enter.

Typing `--help` after a command utility name will provide information about a utility, including the parameters that you can use.

5. What are the operations that you can perform with the `configurebackuplocation` utility?

You will now change the backup location to your shared network folder.

- d Restore File Explorer and browse to `\\EADM\EsriTraining\EADM\ManageArcGISDataStore`.
- e Open the Backups.txt file and copy the command to change the backup directory.
- f In the Command Prompt, paste the command that you copied.
- g Press Enter to execute the command.

You will now place ArcGIS Data Store back in readwrite mode as you are done changing settings.

- h In the Command Prompt, type `changedatastoremode.bat readwrite --store relational`.

Hint: To minimize typing, you can use the up arrow key on your keyboard to restore a previously executed command. You can restore your previous command to place ArcGIS Data Store in readonly mode, and then modify the command to switch ArcGIS Data Store to readwrite mode.

- i Press Enter.
- j When you see the prompt asking if you want to continue, type `yes` and press Enter.

Finally, you will confirm that the backup directory is set correctly.

- k Execute the `describedatastore` utility as you did previously.

```
Information for relational data store ds_yp7f8bym
=====
Backup location.....\\image\RelationalDataStoreBackups
Backup schedule.....{"schedule-starttime":"00:00:00","schedule-frequency":"Every 4 DAYS"}
Days backup retained.....7
Data store status.....Started
SSL enabled.....true
Threshold for READONLY mode.....1024MB
Registered machines.....EBASE.AD.LOCAL
Maximum connections.....150
Owning system URL.....https://ebase.ad.local/server
Portal for ArcGIS URL.....https://ebase.ad.local/portal
Number of connections.....0 connection(s) to managed database
Data store mode.....READWRITE
Is point-in-time recovery enabled...No
```

In this step, you changed the backup location for your relational data store to a shared network folder.

Step 4: Take a manual backup

You will now perform a manual backup of your relational data store.

a In the Command Prompt, type `backupdatastore.bat --store relational` and press Enter.

b When you see the prompt asking if you want to continue, type `yes` and press Enter.



The time that it takes to create the backups and the size of the backups will depend on the size and amount of your data.

c When the operation completes, type `listbackups.bat --store relational` and press Enter.

The output of the `listbackups` utility will show details about the manual backup that you just performed, as well as any other scheduled backups that may have been performed automatically.

d Close all open windows on the EBASE VM.

e Return to the IMAGE VM (orange desktop) and restore File Explorer.

f In File Explorer, view the contents of your shared folder, located at **C:\RelationalDataStoreBackups**.

You will see your backup files in your new folder. Unlike scheduled backups for the relational data store, manual backups do not clear themselves from the disc and must be deleted manually after they are no longer needed.

g Close File Explorer.

In this exercise, you updated the backup location for the relational data store to write to a shared file directory and created a manual backup.

Lesson review

1. When publishing a 3D mesh or a point cloud layer, which ArcGIS Data Store type stores the data?
 - a. Spatiotemporal big data store
 - b. Raster data store
 - c. Tile cache data store
 - d. Object store

2. A relational data store and a tile cache data store can coexist on a single machine, but the spatiotemporal big data store should be installed on its own machine.
 - a. True
 - b. False

3. All hosted services use ArcGIS-managed data, and all services that use ArcGIS-managed data are hosted.
 - a. True
 - b. False

Answers to lesson 7 questions

Types of ArcGIS-managed data (page 7-4)

1. Relational data store
B (Hosted feature layers)
2. Tile cache data store
D (Hosted scene layers)
3. Spatiotemporal big data store
A (High-volume observation data)
4. Raster store
C (Hosted imagery layers)
5. Graph store
F (Knowledge graphs)
6. Object store
E (Cached query responses)

Exercise 7: Manage ArcGIS Data Store backups (page 7-14)

1. What is displayed for the Network Path of this folder?
Not Shared
2. In the RelationalDataStoreBackups Properties dialog box, what is displayed for the Network Path now?
\\IMAGE\RelationalDataStoreBackups
3. Where is the current backup location for the relational data store?
C:\arcgis\arcgisdatastore\backup\relational

Answers to lesson 7 questions (continued)

4. Is point-in-time recovery enabled?

No

5. What are the operations that you can perform with the `configurebackuplocation` utility?

Change, list, register, set default, and unregister

This lesson is the last lesson that focuses on reviewing and comparing the differences between user-managed and ArcGIS-managed data and the impacts of optimization dependent on the storage option. Optimizing services that use ArcGIS-managed data is often very different than optimizing services that use user-managed data. While services that interact with ArcGIS-managed data are typically managed by ArcGIS Enterprise, there are still administrative workflows and strategies that can be used to optimize these services. This lesson will conclude the data management and optimization discussion by focusing on the unique role that ArcGIS-managed data plays in ArcGIS Enterprise.

Topics covered

- Services that use or interact with ArcGIS-managed data
- Optimizing hosted feature services
- Distributed data enabling distributed analysis.

Services that use or interact with ArcGIS-managed data

All services must be able to access data for visualization and use; however, ArcGIS-managed data has a unique role in ArcGIS Enterprise. There are three distinct ways to create web layers and corresponding services that use ArcGIS-managed data: choosing to copy all when publishing from ArcGIS Pro, creating web layers directly in the ArcGIS Enterprise portal, and running analysis in ArcGIS Enterprise.

ArcGIS Pro -Copy all

Similar to services that leverage user-managed data, ArcGIS-managed data provides a storage option to access and manage the underlying data for services. When selecting the copy all option upon publishing in ArcGIS Pro, the data will be copied and managed in ArcGIS Enterprise. The layer and service will leverage the data from an ArcGIS-managed data store, and the type of ArcGIS-managed data store will depend on the type of layer that is published.

That service that uses ArcGIS-managed data may or may not be hosted. Feature layers, tile layers, and vector tile layers published with the copy all data option will create a hosted service. Map image layers and imagery layers that are published by copying the data will create a non-hosted service that uses ArcGIS-managed data.

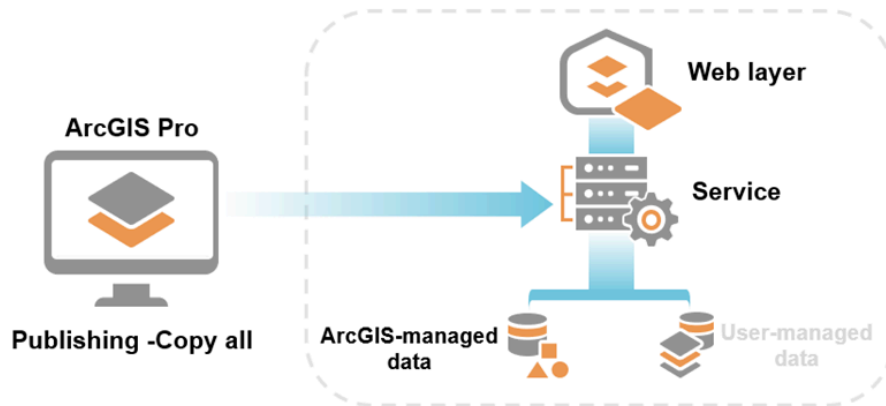


Figure 8.1. When making the decision to copy all data upon publishing from ArcGIS Pro, the layer and corresponding service will use ArcGIS-managed data.

Services that use or interact with ArcGIS-managed data (continued)

ArcGIS Enterprise portal

When web layers are created and published directly in the ArcGIS Enterprise portal, the corresponding service will always be hosted, and the data that it references will always be ArcGIS-managed. For example, uploading and publishing a shapefile in My Content or adding a CSV file containing addresses in Map Viewer will result in a hosted feature layer. ArcGIS-managed data stores are required to take advantage of creating layers directly in ArcGIS Enterprise.

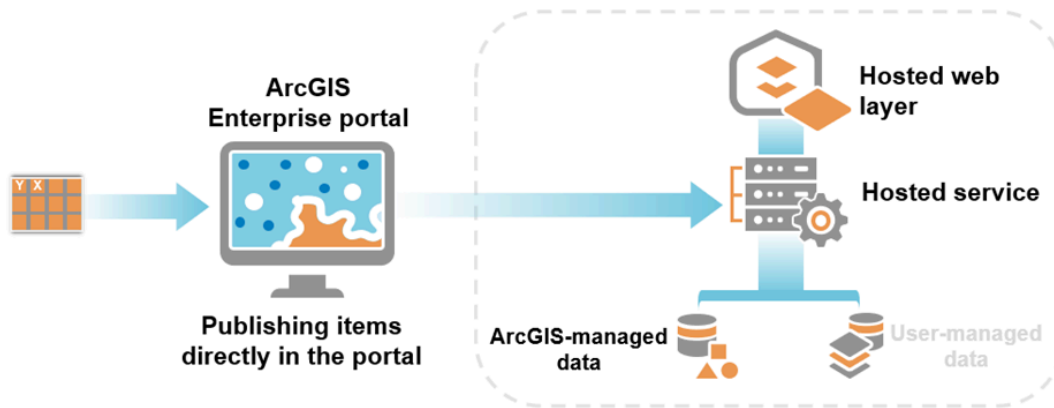


Figure 8.2. Publishing layers directly in ArcGIS Enterprise portal will result in a hosted layer that uses ArcGIS-managed data.

ArcGIS Enterprise analysis

The ArcGIS-managed data storage option is required to conduct analysis within ArcGIS Enterprise. Although the input to analysis may or may not use ArcGIS-managed data, and the geoprocessing services performing the analysis are non-hosted services, the output will always be a hosted service referencing ArcGIS-managed data. ArcGIS-managed data is a pivotal piece of ArcGIS Enterprise as you would not be able to conduct analysis without it.

Services that use or interact with ArcGIS-managed data (continued)

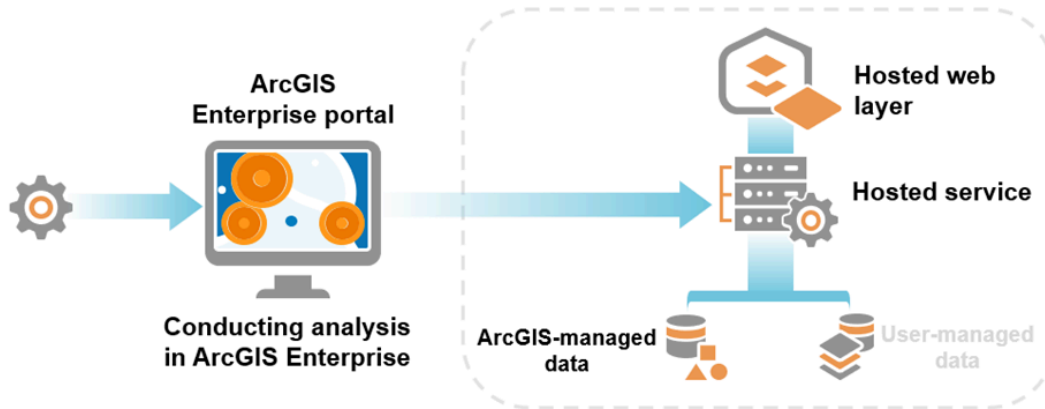


Figure 8.3. Running either the standard feature analysis or raster analysis tools will result in a hosted layer and service that uses ArcGIS-managed data.

Optimizing hosted feature services using the object store

The newest ArcGIS Data Store type, object store, optimizes hosted feature services. Although the data for hosted feature services is stored in the relational data store, the object store caches query responses for hosted feature services.

By saving the query responses, the rendering and performance of the hosted feature service is optimized because the ArcGIS Server accesses the already executed response from the object store, instead of having to execute the query over and over again each time that a new request is made.

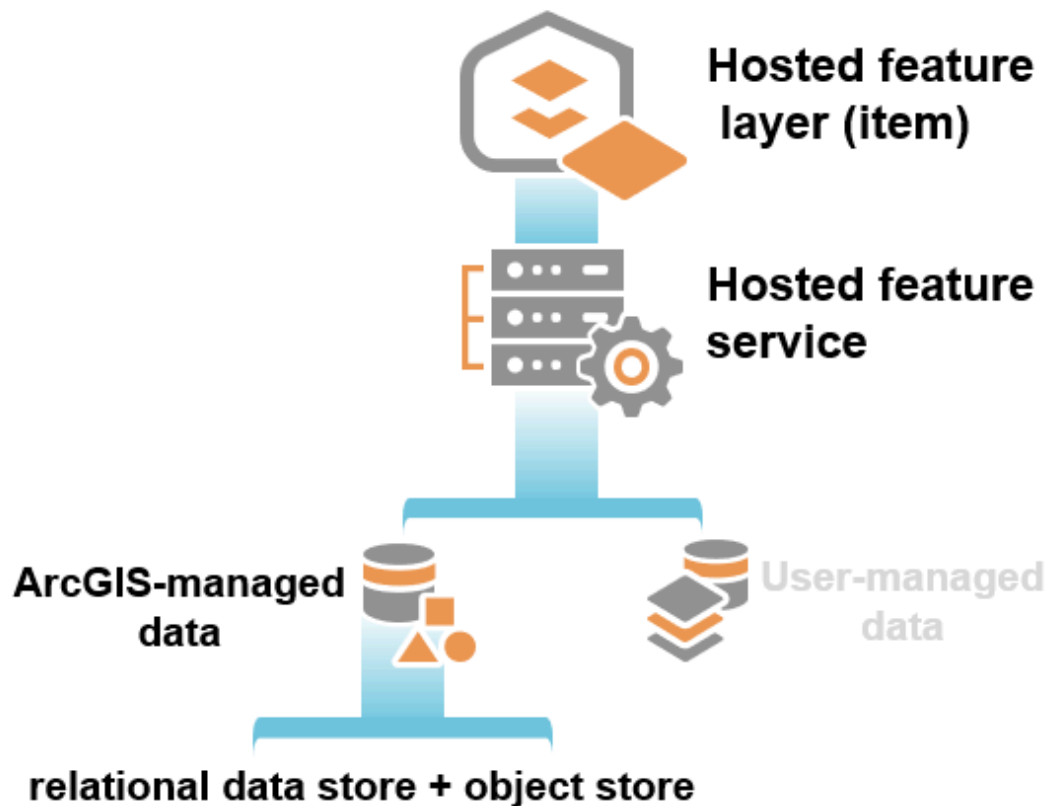


Figure 8.4. The object store can be configured to optimize hosted feature services by storing query responses.

Optimizing hosted feature services using the object store (continued)

Which layers from your organization could benefit from response caching?

Optimizing hosted feature services using the object store video summary

Take advantage of improved drawing times for your hosted feature services

To take advantage of the response caching for your hosted feature services, the object store needs to be configured. The hosted feature services will then need to be configured to leverage the object store. As demonstrated in the video, the ArcGIS Server Admin API is used to run the update definition command for the hosted feature service to enable the layer cache. After it has been configured, the end user should see an improved rendering speed for the hosted feature service, which can be measured with development tools in the web browser.

Distributed data enables distributed analysis

Distributed data is a storage method in which data is stored and accessed in multiple pieces. Storing data in multiple parts optimizes analysis within ArcGIS Enterprise because it enables distributed analysis.

There are several types of ArcGIS-managed data that are stored in a distributed manner, such as the Cloud Raster Format (CRF) and data stored using a multi-node configuration of the spatiotemporal big data store. Storing data in this way serves several roles in analysis workflows:

- **Distributed analysis:** Analysis tasks are completed by multiple processes working in parallel, which is especially useful for very large jobs that are too big for a single process to handle in a reasonable amount of time. These processes can even be distributed across multiple machines. Distributed analysis is inherently scalable; you can scale resources up when needed and scale resources down when they are no longer needed.
- **Concurrent reading from and writing to a single dataset:** Analysis jobs that are divided into multiple parts may be able to read data and write results at the same time, which supports higher data write throughput for large analysis jobs. For example, rasters stored in Cloud Raster Format (CRF) are divided into bundles that can be accessed individually, which enables separate processes to read bundles concurrently and write to different bundles in parallel.

Distributed analysis example

The following graphic shows an example of how these two concepts apply during a distributed raster analysis operation. A Raster Analysis geoprocessing service running on an ArcGIS Image Server site divides the source raster extent by bundle boundaries and assigns raster processing jobs. The raster processing is divided among multiple service instances. Each instance concurrently writes the results to the output dataset, which is stored as a CRF.

Distributed data enables distributed analysis (continued)

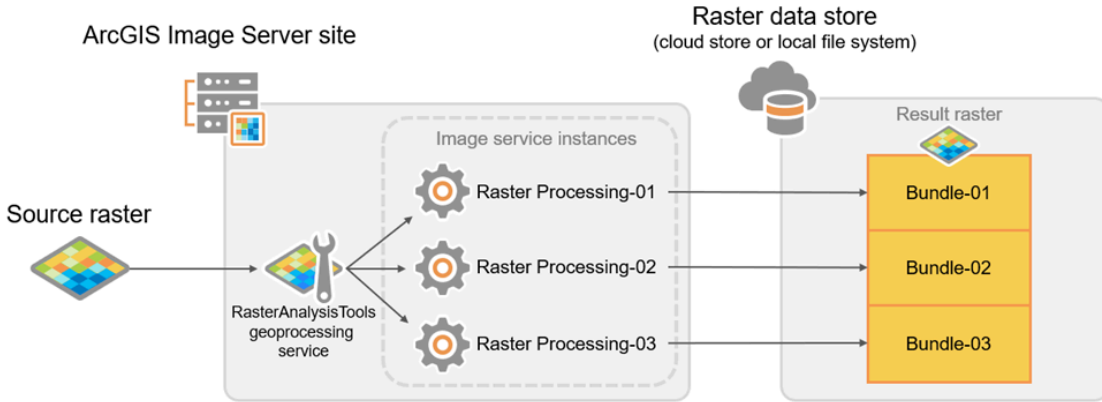


Figure 8.5. The Cloud Raster Format (CRF) facilitates parallel processing by distributing analysis tasks across multiple processes that work concurrently to produce a single result raster.

Exploring distributed raster analytics results

Cloud raster format (CRF) is an example of distributed raster data. When viewing the structure of the data in File Explorer, you can see how the data is divided into levels of bundle files, which enables distributed analysis. Interacting with the data in a client application such as ArcGIS Pro feels very similar to any other layer from a user perspective; the data can be added to a map and used as a parameter in geoprocessing tasks in the same way as other layers in the map.




Figure 8.6. Rasters stored using the CRF format exemplify distributed data.

CRF is well-suited for large rasters because the source raster is divided into multiple directories and files. When accessed in ArcGIS Pro, each required bundle is read and cached locally.

Determining a service optimization strategy

Strategies to optimize services that use ArcGIS-managed data will vary, depending on the workflow and the type of service. Additionally, ArcGIS-managed data is unique in that any time that analysis is run in ArcGIS Enterprise, the output will always be ArcGIS-managed, impacting how you may optimize services that interact with ArcGIS-managed data.

For each of the following examples, determine the strategy to optimize a hosted service given the scenario.

Scenario 1: Raster analysis

An image analyst is attempting to conduct raster analysis on a mosaic dataset containing hundreds of images in ArcGIS Pro. Because they are running the analysis on their local machine, they are consuming most of their local machine's system resources, preventing them from completing their other daily workflows while the analysis is running. They are contemplating moving the analysis to ArcGIS Enterprise but are still concerned about processing time and potentially missing deadlines.

- 1. What aspect of raster storage can alleviate concerns about processing imagery server side?**

Determining a service optimization strategy (continued)

Scenario 2: Numerous query requests

A real estate analyst is receiving numerous query requests on their hosted feature services and observing a reduced performance.

2. What strategy should be deploying to optimize the hosted feature services?

Exercise 8

Use ArcGIS Server to optimize feature analysis

Imagine that a new project recently kicked off, requiring analysts within the department to perform additional analysis to meet project deadlines each week. With the uptick of analysis, the analysts have begun to see a degradation of performance in ArcGIS Enterprise. As the GIS administrator, you were notified of the poor performance, and you have begun reviewing logs and statistic reports. You find that the SpatialAnalysisTools system service is receiving timeouts. There are more requests being made to this service due to the work being done on the new project. You will scale the SpatialAnalysisTools system service to meet the increased demand.

In this exercise, you will perform the following tasks:

- Scale the SpatialAnalysisTools system service.
- Run standard feature analysis and inspect the hosted feature layer output.
- Explore deletion behavior for ArcGIS-managed data.

Step 1: Inspect SpatialAnalysisTools service instances

You will first inspect the SpatialAnalysisTools service instances. This system geoprocessing service is used to handle requests that are made when running the standard feature analysis tools. The hosting server site runs the system geoprocessing services that complete the analysis and will also host the service output from the analysis.

- a Connect to the EBASE VM (blue desktop) using the Remote Desktop Connection Manager.
- b Open a web browser and click the EBASE-Server Manager bookmark (<https://ebase.ad.local/server/manager>).
- c Sign in with your administrative credentials (**portaladmin** username and **Esri.4.GIS** password), if necessary.
- d Under Folders, click the System folder to view the system services for the hosting server site.
- e In the search field, type **Spatial**.



SpatialAnalysisTools is a system geoprocessing service. Geoprocessing services require dedicated service instances and, therefore, have the ability to configure Minimum and Maximum instances for the service. By default, the total number of instances running is one; however, the service has the capacity to spin up one more additional service instance if the demand requires it.

- f On the EBASE VM, right-click the Windows Taskbar and choose Task Manager.
- g In Task Manager, locate the SpatialAnalysisTools ArcSOCs processes.

Hint: In Task Manager, click the Details tab, and then scroll to find the SpatialAnalysisTools ArcSOCs.

1. How many SpatialAnalysisTools ArcSOCs processes are running on the EBASE VM?

```
+UseParallelGC -Dservice=System.SpatialAnalysisTools.GPServer --add-open
+UseParallelGC -Dservice=System.SpatialAnalysisTools.GPServerSync --add-i
```

There are two ArcSOCs that are available in Task Manager for the SpatialAnalysisTools service: GPServerSync and GPServer. The GPServer manages the queue of incoming requests, and the GPServerSync performs the actual geoprocessing task. When increasing the max number of service instances, you are increasing the number of GPServerSync ArcSOCs that can spin up so you can handle more simultaneous geoprocessing tasks.

h Review the system resources consumed by the SpatialAnalysisTools ArcSOCs.

2. Which system resources are consumed by the SpatialAnalysisTools?

Although analysis is not occurring at this time, you have reviewed statistics reports and found that the SpatialAnalysisTools service has received timeouts throughout the past few days. You will scale the system service to meet the demand.

i Minimize the Task Manager window to review later.

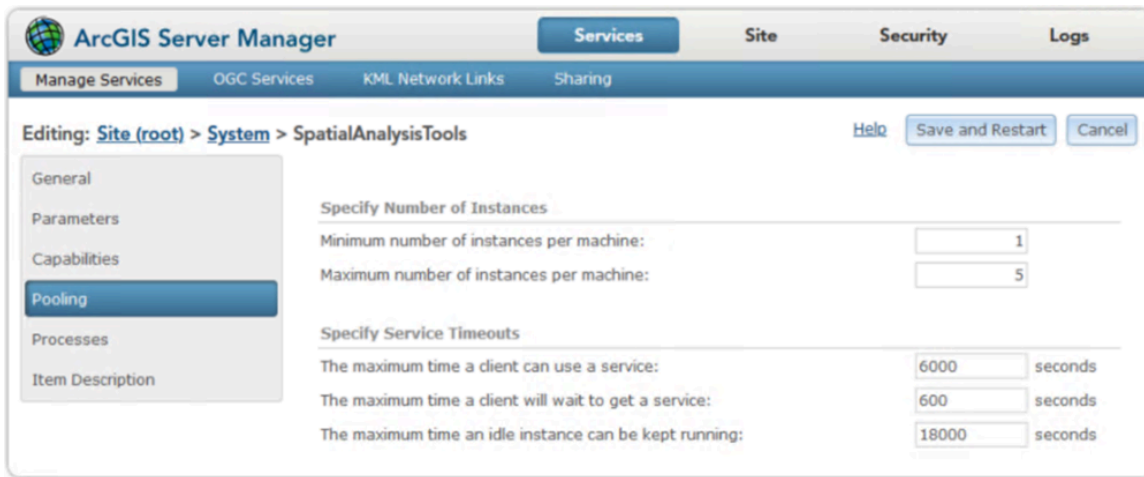
Step 2: Increase instances for the SpatialAnalysisTools service

In this step, you will scale the SpatialAnalysisTools system service by increasing the number of maximum service instances in ArcGIS Server Manager for the hosting server.

a In ArcGIS Server Manager, click the SpatialAnalysisTools service to change the service properties.

b On the left, click the Pooling tab.

- c Under Specify Number Of Instances, for Maximum Number Of Instances Per Machine, type 5.



- d Click Save And Restart.
- e Under Folders, click the System folder, if necessary, and then search for the SpatialAnalysisTools service again.
- f Notice that Maximum Instances has been increased to 5 but only 1 instance is currently running.



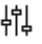
3. How many ArcSOCs should be available in Task Manager after increasing the maximum number of instances for the service?

You increased the number of services instances; therefore, you increased capacity of the system service to handle additional demand.

- g Minimize the Task Manager window.
- h Leave the ArcGIS Server Manager web browser tab open.

Step 3: Set up the analysis

In this step, you will set up the analysis to observe the SpatialAnalysisTools instances respond to requests.

- a Open a new web browser tab and click the ArcGIS Enterprise bookmark (<https://ebase.ad.local/portal/home>).
- b Sign in with your administrative credentials (**portaladmin** username and **Esri.4.GIS** password), if necessary.
- c Click Content, and then click the My Organization tab.
- d In the Search My Organization field, type **Tree Survey**.
- e For the Tree Survey feature layer, click the More Options button  and choose Open In Map Viewer.
- f On the right side of the page, from the Settings toolbar, click the Analysis button , and then click Tools.
- g Click the Filter button  and choose Feature Analysis.
- h Expand Analyze Patterns, and then click Calculate Density.
- i For Input Point Or Line Features, click Layer and choose Tree_Survey.
- j Under Result Layer, for Output Polygon Name, type **Tree Survey Density Test 1**.
- k Leave all other default parameters.
- l Do not click Run yet.

You will repeat the analysis setup to run two analysis tasks simultaneously.

- m Open an additional web browser tab and click the ArcGIS Enterprise bookmark (<https://ebase.ad.local/portal/home>).
- n Sign in with your administrative credentials (**portaladmin** username and **Esri.4.GIS** password), if necessary.

- o From the Content page, on the My Organization tab, search for the Tree Survey feature layer and open the layer in Map Viewer.
- p Open the Calculate Density analysis tool.
- q For Input Point Or Line Features, click Layer and choose Tree_Survey.
- r Under Result Layer, for Output Polygon Name, type **Tree Survey Density Test 2**.
- s Leave all other default parameters.

Step 4: Run the analysis

Now that the two analysis tasks are set up, you will run both analysis tasks simultaneously to review how the service instances respond.

- a In Map Viewer, click Run.
- b Return to the previous web browser tab with Map Viewer and click Run.
- c Restore Task Manager and view the SpatialAnalysisTools ArcSOCs.

4. How many ArcSOCs are now available for SpatialAnalysisTools.GPServerSync?

```
-Dservice=System.SpatialAnalysisTools.GPServer --add-opens=java.base/java.net=ALL...
-Dservice=System.SpatialAnalysisTools.GPServer --add-opens=java.base/java.net=ALL...
-Dservice=System.SpatialAnalysisTools.GPServerSync --add-opens=java.base/java.net...
-Dservice=System.SpatialAnalysisTools.GPServerSync --add-opens=java.base/java.net...
```


Two additional ArcSOCs were spun up to work on an additional request.

- d Open the ArcGIS Server Manager web browser tab.
- e Notice the number of instances running and in use for SpatialAnalysisTools.
- f After reviewing ArcGIS Server Manager and Task Manager, wait for the analysis to be completed.



It may take several minutes for the analysis to be completed.

You will know when the analysis is complete when Instances In Use returns to 0.



SpatialAnalysisTools (Geoprocessing Service)
 The SpatialAnalysisTools service is used for supporting S
 ArcGIS.

| | |
|--------------------|---------|
| Status: | Started |
| Instances Running: | 3 |
| Instances in Use: | 0 |
| Maximum Instances: | 5 |

Analysis outputs are created as ArcGIS-managed hosted feature layers. You were able to optimize the SpatialAnalysisTools geoprocessing service that writes the ArcGIS-managed data output to the relational data store.

Step 5: Inspect hosted feature layer output

In this step, you will review the hosted feature layers and corresponding services output from the analysis tasks that you ran in the previous step.

- a Restore the web browser tab that contains the map with the analysis result Tree Survey Density Test 1.
- b Click the Menu button and choose Content.
- c Accept the option to leave the map without saving changes.
- d Click My Organization, and then search for **Tree Survey**.
- e Notice that the analysis outputs are hosted feature layers.

Similar to what you discovered when publishing hosted services from ArcGIS Pro, there are numerous capabilities that are available for the layer in the ArcGIS Enterprise portal.

- f Restore the ArcGIS Server Manager web browser tab.
- g On the left, under Folders, click the Hosted folder.

Both Tree Survey Density Test services are listed as a service on the hosting server.

- h Click Tree Survey Density Test 1.
- i Notice that there are no service properties to configure as it is a hosted service.

As this action was a test to review the scaling potential of the service, you do not need to maintain the output of the analysis. You will now delete the new hosted feature layers that were

created from the analysis.

Step 6: Explore deletion behavior of a hosted feature layer

Now you will explore deletion behavior of hosted data as it relates to layers and services.

- a Restore the browser tab with the content in My Organization displayed.
- b Check the boxes for the Tree Survey Density Test 1 and Tree Survey Density Test 2 feature layers.
- c Click Delete, and then click Delete to confirm the deletion.
- d Restore the ArcGIS Server Manager web browser tab.
- e Click the Services tab at the top of the page to return to the list of services, and then under Folders, click the Hosted folder.

The Tree Survey Density Test services no longer exist. When you delete a portal item referencing a hosted layer, the service and the data are deleted.

- f Close the web browser without saving any changes to the maps.

Due to the SpatialAnalysisTools system service receiving timeouts, you determined that it was necessary to scale the SpatialAnalysisTools system service to meet the increased demand. After increasing the maximum number of service instances, you ran multiple analysis tasks simultaneously to review how the service instances responded to the requests.

After the test, you deleted the hosted feature layer outputs, which subsequently deleted the layers, services, and underlying data. Analysts will have additional system resources available for their analysis tasks.

Lesson review

1. When performing analysis in ArcGIS Enterprise, which type of service is always created?
 - a. Non-hosted service
 - b. Hosted service
 - c. Feature service
 - d. Map service

2. How does the object store optimize hosted feature services?
 - a. Creates pre-rendered tiles
 - b. Provides additional storage
 - c. Caches query responses
 - d. Distributes the data into bundles

3. How does the role of distributed data impact scalable distributed analysis?

Answers to lesson 8 questions

Determining a service optimization strategy (page 8-11)

Scenario 1: Raster analysis

1. What aspect of raster storage can alleviate concerns about processing imagery server side?

The output of raster analysis in ArcGIS Enterprise is divided into multiple bundles files, which provide the ability for multiple processing and analysis jobs to run and write results more efficiently. Because the distributed data provides the ability for the distributed analysis, the raster analytics is able to be completed within a few hours, rather than a few days. Additionally, due to the analysis taking place server side, it does not impact system resources on the image analyst's local machine; therefore, they have the ability to use their local machine as necessary while waiting for the analysis to complete.

Scenario 2: Numerous query requests

2. What strategy should be deploying to optimize the hosted feature services?

Deploy and leverage the object store for faster drawing of the hosted feature services.

Exercise 8: Use ArcGIS Server to optimize feature analysis (page 8-13)

1. How many SpatialAnalysisTools ArcSOCs processes are running on the EBASE VM?

There are two ArcSOCs that are currently running for the SpatialAnalysisTools service.

2. Which system resources are consumed by the SpatialAnalysisTools?

Although the ArcSOCs are not consuming any of the machine's CPU and are sitting idle, the ArcSOCs are still consuming Memory or RAM.

Answers to lesson 8 questions (continued)

3. How many ArcSOCs should be available in Task Manager after increasing the maximum number of instances for the service?

There are two ArcSOCs that are available in Task Manager: one GPServerSync and one GPServer. However, now the service has the capacity to increase the number of GPServerSync ArcSOCs to five, if needed.

4. How many ArcSOCs are now available for SpatialAnalysisTools.GPServerSync?

Two

The spatiotemporal big data store is one of the types of data stores that is created by ArcGIS Data Store. The spatiotemporal big data store is necessary to persist real-time observational data and is designed to support high-velocity write throughput.

The spatiotemporal big data store works closely with two server licensing roles. ArcGIS GeoAnalytics Server requires a spatiotemporal big data store to store its analysis results and can use data stored there for analysis input. ArcGIS GeoEvent Server can stream or consume real-time event-based data, which can be archived in the spatiotemporal big data store.

Topics covered

Understanding the spatiotemporal big data store

Persisting real-time data in ArcGIS Enterprise

Supporting distributed vector analysis

Understanding the spatiotemporal big data store

Earlier in the course, you learned about ArcGIS Data Store types and their relationship with data management strategies. This lesson focuses on the spatiotemporal big data store and its role in an ArcGIS Enterprise deployment.

The spatiotemporal big data store serves several overarching roles in an ArcGIS Enterprise deployment.



Figure 9.1. The spatiotemporal big data store provides storage using a multi-node architecture, visualization capabilities, and support for distributed analysis.

- **Storage:** The spatiotemporal big data store provides ArcGIS Enterprise with a mechanism for persisting huge amounts of high-volume observational data. Various apps such as ArcGIS Field Maps and ArcGIS Mission can write to the spatiotemporal big data store, which uses a multi-node architecture that is inherently scalable and resilient.
- **Visualization of big data:** Data stored in the spatiotemporal big data store can be visualized using map and feature services that you can work with in the Enterprise portal. These map and feature services are optimized for big data and include features such as aggregating your data based on scale to help interpret very large datasets.
- **Distributed analysis:** After vector or tabular data is stored, the data can be used for analysis. Big data analysis tools, which can be made available in Map Viewer Classic and ArcGIS Pro, use distributed analysis capabilities, speeding up analyses of very large datasets. To perform distributed analysis, these tools use parallel processing to quickly analyze large amounts of data. When these tools are run, the results can be stored in the spatiotemporal big data store.

Understanding the spatiotemporal big data store (continued)

Integrating the spatiotemporal big data store with ArcGIS Enterprise

The spatiotemporal big data store is closely integrated with two server licensing roles:

- ArcGIS GeoEvent Server enables real-time analytics.
- ArcGIS GeoAnalytics Server provides big data processing and analysis capabilities.




Figure 9.2. The spatiotemporal big data store serves as the data storage solution that supports distributed analysis and persisting real-time observational data in ArcGIS Enterprise.

After data is ingested and analyzed using GeoEvent Server from real-time, high-volume event data from sources like the Internet of Things (IoT), the spatiotemporal big data store can persist the event data. Thus, GeoEvent Server can write data to the spatiotemporal big data store. GeoAnalytics Server, however, can use data stored in the spatiotemporal big data store as the input for analysis, in addition to writing analysis results to the spatiotemporal big data store.

Map and feature services can access data in the spatiotemporal big data store and are treated as another type of portal item that users can discover, visualize, and share.

Understanding the spatiotemporal big data store (continued)



It is important that administrators monitor the disk usage closely because the archived data can accumulate quickly.



ArcGIS Enterprise Help: *A quick tour of GeoEvent Server*

ArcGIS Enterprise and the spatiotemporal big data store

After the spatiotemporal big data store is installed, it is registered with the ArcGIS Enterprise hosting server.

Architecturally, components such as a GeoEvent Server site or GeoAnalytics Server site access the spatiotemporal big data store by first making a connection to the ArcGIS Enterprise hosting server.

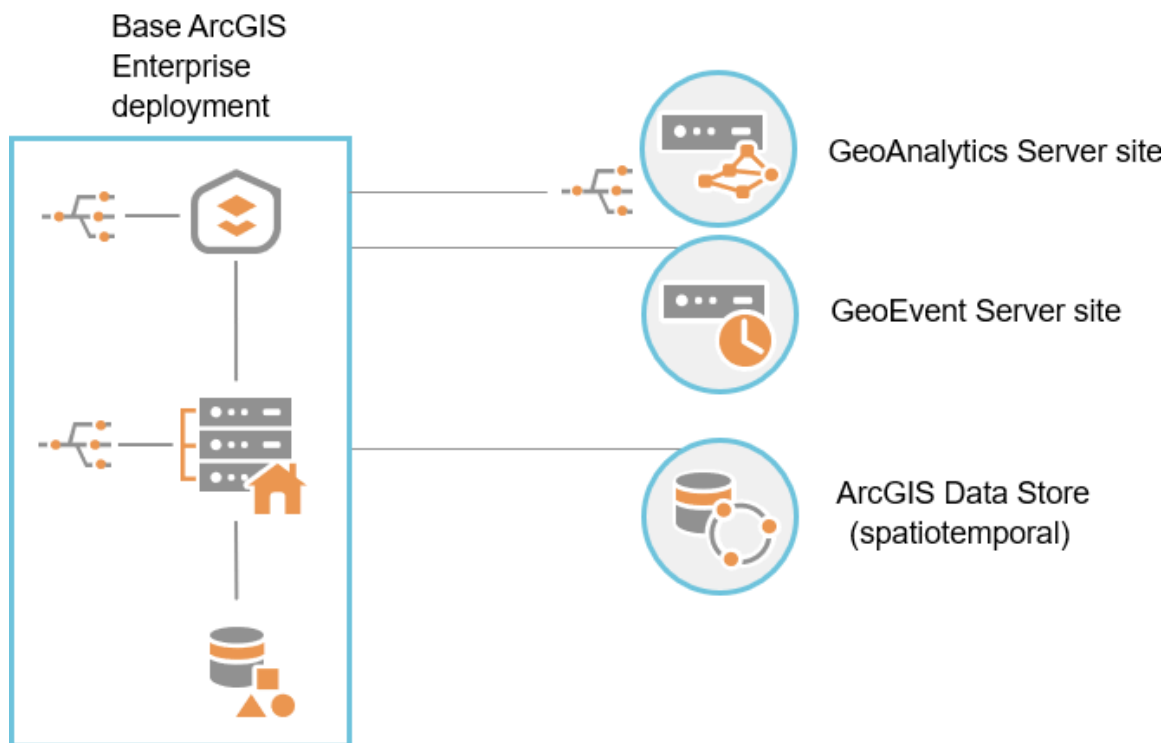


Figure 9.3. In this example, a spatiotemporal big data store is configured with a single node on a dedicated machine. It is registered with the hosting server of the base ArcGIS Enterprise deployment to make the data store available to other ArcGIS Enterprise components such as a GeoAnalytics Server site and GeoEvent Server site.

Accessing data in the spatiotemporal big data store

In addition to GeoAnalytics Server and GeoEvent Server being able to write to the spatiotemporal big data store, apps such as ArcGIS Indoors, ArcGIS Mission, and ArcGIS Field Maps can store tracks and other location information using the spatiotemporal big data store. The hosted feature layers that reference data in the spatiotemporal big data store have the same

ArcGIS Enterprise and the spatiotemporal big data store (continued)

editing functionality as layers that expose data in the relational data store.

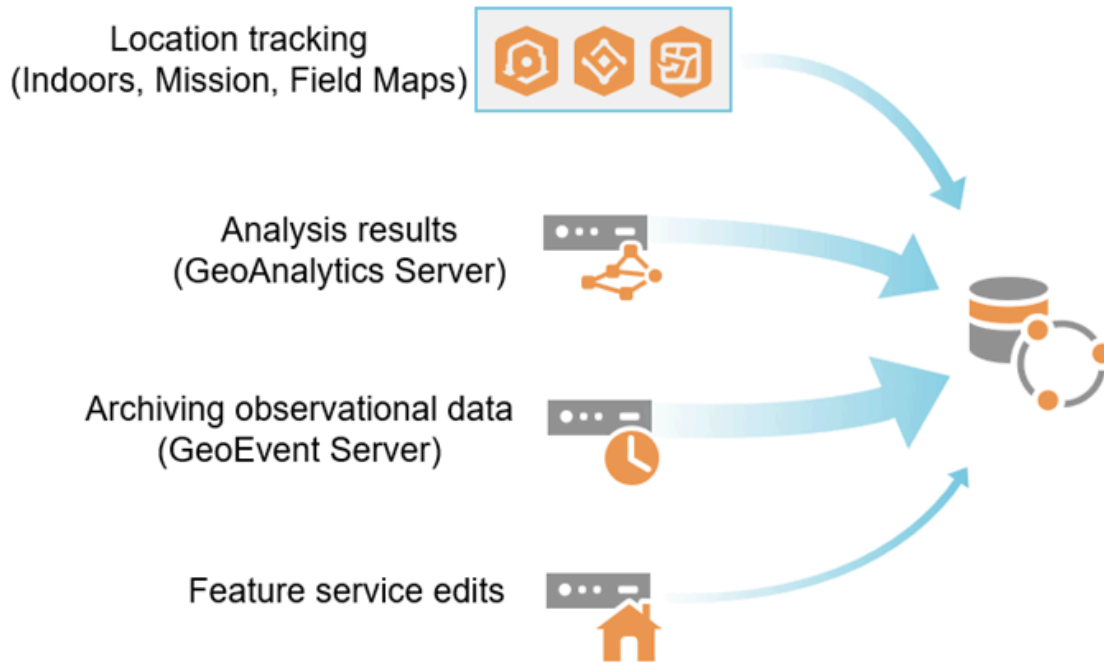


Figure 9.4. This example illustrates various components that write data to the spatiotemporal big data store. In this graphic, the size of the arrows indicates the relative amount of data that is typically sent to the spatiotemporal big data store.

Apps such as ArcGIS Indoors, ArcGIS Mission, and ArcGIS Field Maps can read tracks and other location information that they store in the spatiotemporal big data store. GeoAnalytics Server can use data stored in the spatiotemporal big data store as input into distributed vector analysis. Hosted map and feature services that expose data within the spatiotemporal big data store can be consumed by client applications such as Map Viewer and ArcGIS Pro.

ArcGIS Enterprise and the spatiotemporal big data store (continued)

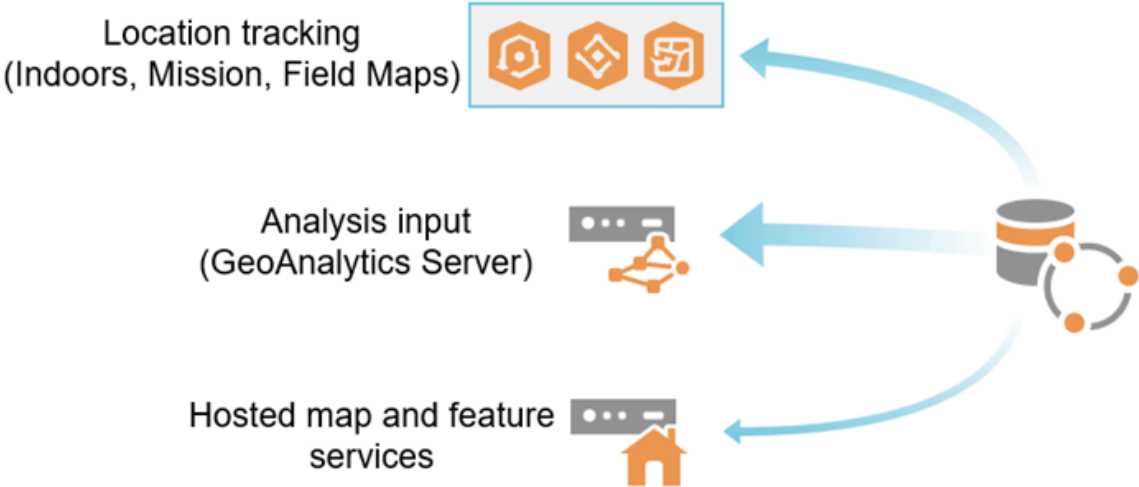


Figure 9.5. This example illustrates various components that read data from the spatiotemporal big data store. In this graphic, the size of the arrows indicates the relative size of typical read operations.

How the spatiotemporal big data store relates to other ArcGIS Enterprise components

The spatiotemporal big data store has different relationships with different components of ArcGIS Enterprise.

Instructions

In the following table, the column on the left describes some examples of ArcGIS Enterprise functions that relate to the spatiotemporal big data store. In the column on the right, you will provide the letter that corresponds with one of the following options:

- A = Dedicated host (no other ArcGIS components installed)
- B = Enterprise hosting server
- C = ArcGIS GeoEvent Server
- D = ArcGIS GeoAnalytics Server

| Functions related to the spatiotemporal big data store | Match letter |
|--|--------------|
| The spatiotemporal big data store is registered with it. | |
| It writes real-time observational data to the spatiotemporal big data store. | |
| Additional spatiotemporal big data store nodes should be installed on it. | |
| It can analyze data from the spatiotemporal big data store. | |

Exercise 9

Explore data in the spatiotemporal big data store

Imagine that you are a transportation analyst who has access to a live stream of bus locations for the city of Los Angeles, California. You want to be able to display live bus locations, but you also want to be able to generalize the data so that you can see where the greatest concentration of buses are located at any given time of day. In addition, you want to be able to filter the data so that you can see specific buses that pertain to specific routes.


In this exercise, you will perform the following tasks:

- Investigate a single-node configuration of the spatiotemporal big data store.
- Use and filter spatiotemporal big data store data.

Step 1: Investigate the spatiotemporal big data store

You will explore a GeoEvent Server configuration on the GEO VM, as well as a spatiotemporal big data store installed on the STBDS VM. Additionally, about 100,000 bus locations have been preloaded into the spatiotemporal big data store, representing a recorded stream of Los Angeles bus locations over a period of time.

- a Connect to the STBDS VM using the Remote Desktop Connection Manager.
- b Sign in to the STBDS VM (brown desktop) as **adAdmin** with the **Esri.4.GIS** password, if necessary.
- c Click the Windows Start icon, and then expand the ArcGIS folder.



- d Click the ArcGIS Data Store Configuration Wizard.
- e In the URL, replace *localhost* with **stbds.ad.local** so that the full URL is `https://stbds.ad.local:2443/arcgis/datastore`, and then press Enter.

Configuration Status

[Help](#)

The following component(s) of ArcGIS Data Store are configured with ArcGIS Server site(s).

- Spatiotemporal

You can now use the Portal URL below to publish and access services.

Portal URL: <https://ebase.ad.local/portal>

[Configure Additional Data Stores](#)

A web browser opens to the ArcGIS Data Store Configuration page.

The only type of ArcGIS Data Store on this machine is the spatiotemporal big data store. The spatiotemporal big data store has been configured with the hosting server, as required.

Step 2: View a spatiotemporal big data store node

Scaling of the spatiotemporal big data store simply requires installing ArcGIS Data Store configured as a spatiotemporal big data store on additional machines and configuring each with the hosting server. It is a best practice to install an odd number of spatiotemporal nodes, each on dedicated hardware, for optimal performance.

- a In the web browser, click the EBASE-Server Admin bookmark (<https://ebase.ad.local/server/admin>) to open the ArcGIS Server Administrator Directory.
- b Sign in with the **portaladmin** username and **Esri.4.GIS** password.
- c Click Data, click Items, and then click NosqlDatabases.

Data Item (/nosqlDatabases)

Data Item Properties

Item Path: [/nosqlDatabases](#)

Item Type: datadir

ID: 06f1df11-b1a3-4d7f-af4b-11d61dbc98bd

Client Path: N/A

Child Items

- [/nosqlDatabases/AGSDataStore_bigdata_bds_qmsbahtc](#)
- [/nosqlDatabases/AGSDataStore_nosqldb_tcs_zactrmgm](#)
- [/nosqlDatabases/AGSDataStore_queue_xu7h8wt](#)

1. Which of these data stores do you believe is the spatiotemporal big data store?

There are two other NoSQL Databases listed: the `_tcs` data store, which is the tile cache data store, and the `_queue`, which is used to support webhooks.

- d Click the link with `_bigdata_bds` in it, and then click the Machines link.

Data Store Machines

Machines:

- [STBDS.AD.LOCAL](#)

Multiple nodes may be necessary to support the level of availability and performance that will meet your business needs. For high availability, use an odd number of nodes; a minimum of at least three nodes is recommended, but five nodes are optimal if there are sufficient hardware resources.

- e Click the STBDS.AD.LOCAL node of the spatiotemporal big data store.

2. Which port must be open for spatiotemporal BDS communication?

Hint: Look for the Database port property.

Internal communication between spatiotemporal big data store machines happens through this port.

Step 3: Explore the spatiotemporal big data store in GeoEvent Manager

Next, you will explore how the spatiotemporal big data store is managed using GeoEvent Manager.

- a Open a new web browser tab and click the GEO-GeoEvent Manager bookmark (<https://geo.ad.local:6143/geoevent/manager>).



If ArcGIS GeoEvent Manager does not appear, use Task Manager to stop the following services on the GEO machine in the following order: ArcGIS GeoEvent Server service, ArcGIS GeoEvent Gateway service, and ArcGIS Server service. Then, restart the ArcGIS GeoEvent Server service. Restarting this service will trigger the other two services to restart.

ArcGIS GeoEvent Manager will open in a new web browser tab and is operating on the geo.ad.local machine.


b Sign in to GeoEvent Manager with the following credentials:



- Username: **siteadmin**
- Password: **Esri.4.GIS**

c Click the Site tab, and then on the left, click the Data Stores tab.

d Under Registered Server Connections, notice the entry for ArcGIS Enterprise On EBASE.

If you recall, the spatiotemporal big data store is registered with your hosting server. By registering a connection to your ArcGIS Enterprise deployment here on the Data Stores page, GeoEvent Server is able to discover and use the spatiotemporal big data store. This connection has already been registered for you, but you will investigate the connection properties.

e For the ArcGIS Enterprise On EBASE connection, click the Edit button , as indicated in the following graphic.

| Registered server connections | | Register server connection | | | |
|-------------------------------|--|----------------------------|---|---|---|
| Status | Name | Type | Sync | Edit | Delete |
| ✓ | ArcGIS Enterprise on EBASE | ArcGIS Enterprise |  |  |  |

By looking at the connection, you can see that this spatiotemporal big data store is registered to ArcGIS Enterprise on the EBASE machine.

3. What are the three types of server connections that can be made?

4. Which URL is used to connect to ArcGIS Enterprise on EBASE?


f Click Cancel to close the Edit Server Connection dialog box.

g On the left, click the Spatiotemporal Big Data Stores tab.

| Data Source | Map Services | Feature Services | Action |
|--|--|--|---|
| LA_Bus_BDS_Write     | LA_Bus_BDS_Write    | LA_Bus_BDS_Write    |  |



If the content does not match the preceding graphic, perform the following steps:

1. Click the Data Stores tab, and then for ArcGIS Enterprise On EBASE, click the Edit button .
2. To cause a disconnect, add an **X** to the end of the username portaladmin.
3. Click Register (the connection will fail).
4. Change the name back to **portaladmin**, and then click Register again to force a new connection.
5. Refresh the web browser, and then click the Spatiotemporal Big Data Stores tab.
6. Wait up to 5 minutes for the information to appear and avoid refreshing the web browser more than once.

In GeoEvent Manager, the data in the spatiotemporal big data store can be accessed and managed through the data source, as well as its optional map service and feature service.

- The **Data Source** column lists the data sources or database containers where the data is stored. You can update data properties such as replication factor, data retention, and max record count. You can also delete all or a subset of the data based on a SQL query.
- The **Map Services** and **Feature Services** columns list the enhanced hosted map and feature services that access the data from a particular data source, within spatiotemporal big data store.

Step 4: View data stored in the spatiotemporal big data store

Next, you will use these services to view the data stored in the spatiotemporal big data store.


- a Under Map Services, click the link for LA_Bus_BDS_Write.
- b Sign in as **portaladmin** with the **Esri.4.GIS** password, if necessary.

The ArcGIS REST Services Directory opens.

5. If you wanted to view this hosted map service in ArcGIS Server Manager, which ArcGIS Server site would you use?

Because these services are hosted, corresponding layer items in the Enterprise portal were created when the services were published. You will explore these layer items next.

- c Open a new web browser tab and click the ArcGIS Enterprise bookmark (<https://ebase.ad.local/portal/home/>).
- d Sign in as **portaladmin** with the **Esri.4.GIS** password, if necessary.
- e Click Content.
- f In the list of content, locate the LA_Bus_BDS_Write map image layer and feature layer.



The LA_Bus_BDS_Write map image layer and feature layer may be on a subsequent page.

The feature layer icon indicates that the data for this layer originates from a spatiotemporal big data store.

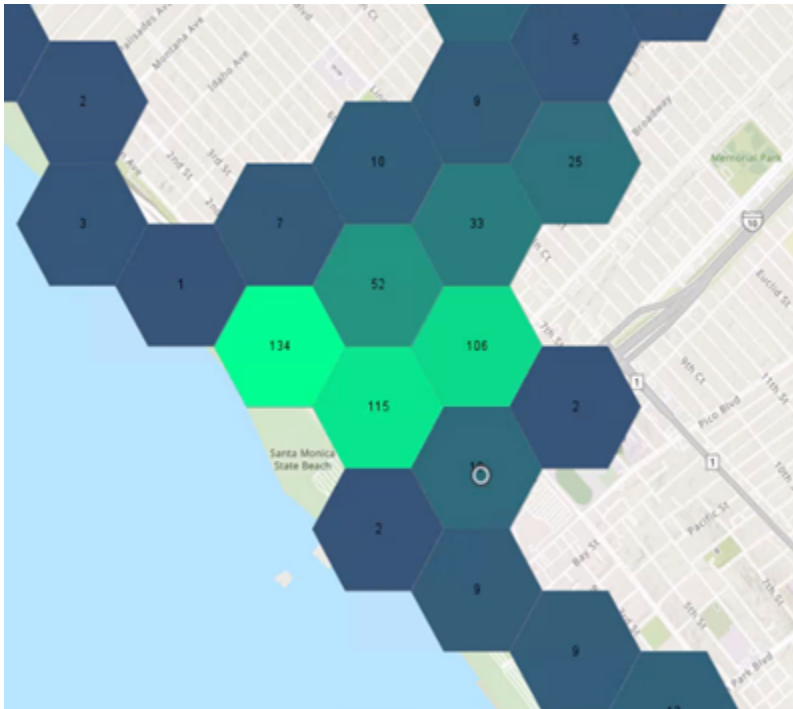
- g To the right of the LA_Bus_BDS_Write map image layer, click the More Options button **⋮** and choose Open In Map Viewer.




Be sure to add the map image layer to the map and not the feature layer.

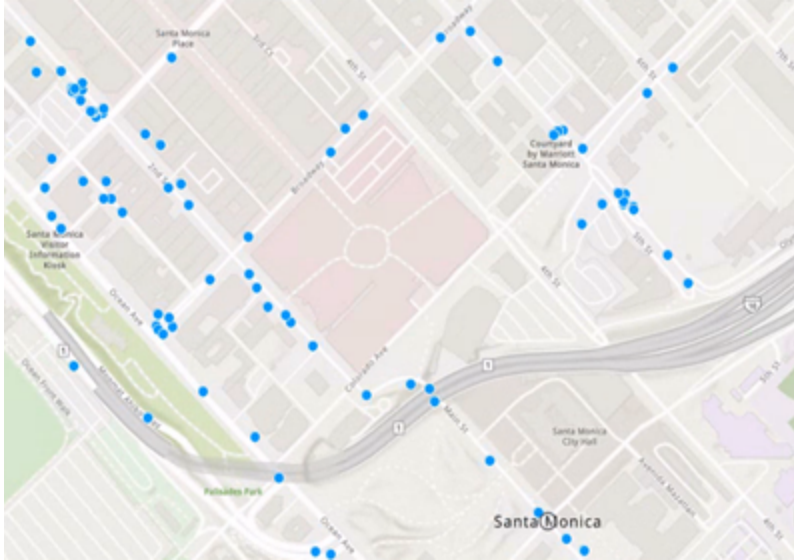
The map image layer is added to Map Viewer in the same way as other layers in the Enterprise portal. The spatiotemporal big data store enables special characteristics for this layer, such as on-the-fly aggregation, which you will explore next.

- h Near the bottom-right corner of the map, click the Search button 🔍.
- i In the search field, type **Santa Monica, CA**, and then press Enter.
- j In the Search Result pop-up window, click Zoom To.
- k Close the Search Result pop-up window.
- l Notice the label and symbology for feature counts contained in each hexagonal aggregation box.





 Your results may vary slightly from the preceding graphic.

- m Zoom in further until the display changes to blue points.




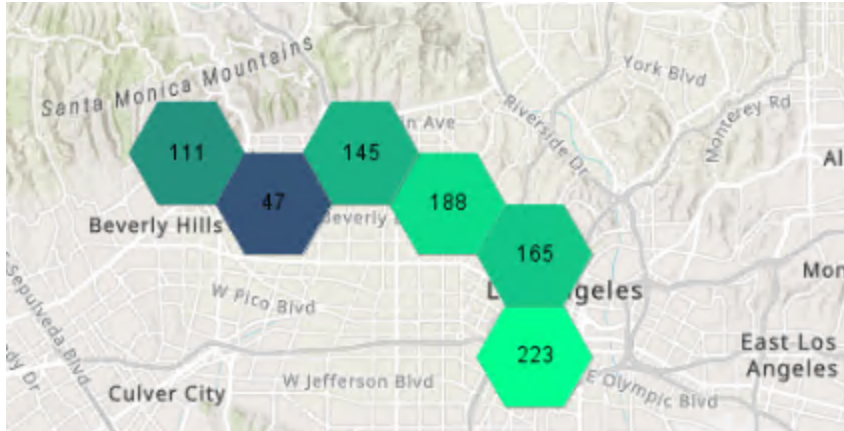
This layer can also be queried and filtered. You will write a query expression to limit the number of features visible in the layer.

- n In the Layers pane, for LA_Bus_BDS_Write, click the Expand button .
- o Click the nested LA_Bus_BDS_Write layer to select it.
- p On the Settings toolbar, click the Filter button  and then click Add Expression.
- q Under Expression, create the following expression: Route_id is **10**.

Expression ***

| | |
|----------|---|
| route_id | ▼ |
| is | ▼ |
| 10 | ▼ |

- r Click Save.
- s Click the Search button  and search for **Los Angeles, CA** to zoom out.
- t In the Search Result pop-up window, click Zoom To.
- u Close the Search Result pop-up window.



Your results may vary slightly from the preceding graphic.



The filtered aggregated data for the buses appears.

Step 5: Manipulate data in a spatiotemporal big data store


Next, you will use ArcGIS Enterprise tools to add and delete data from the spatiotemporal big data store.


- a Remove the LA_Bus_BDS_Write map image layer from the map.

Hint: In the Layers pane, for the parent layer, click the Options button ******* and choose Remove.

- b Click Add.
- c Locate the LA_Bus_BDS_Write feature layer and click Add.
- d To the left of Add Layer, click the Back button **<**.
- e On the Settings toolbar, click the Edit button .
- f In the Editor pane, under Create Features, click LA_Bus_BDS_Write.
- g Click anywhere in the map to add a new feature at that location.
- h Click Create without entering any attribute values.
- i Click the Edit button  to stop editing.


Now you will use GeoEvent Manager to selectively delete some of the data from the original spatiotemporal big data store data source for the LA_Bus_BDS_Write feature service.

- j Restore the web browser tab that shows the Spatiotemporal Big Data Stores tab in ArcGIS GeoEvent Manager.
- k To the right of the LA_Bus_BDS_Write data source, click the Delete Data button , as indicated in the following graphic.



- l Next to Delete Features, select By Where Clause.
- m In the Where field, type the SQL expression `route_id <> 10` to delete all points except for those points corresponding to bus route 10.

You are about to delete features from the Data Source **LA_Bus_BDS_Write**. Please select 'All', or specify a selective 'By Where Clause'. ('By Where Clause' example that deletes all features: `1=1`)

Delete Features:  All By Where Clause

Where: 


- n Click Delete.



If prompted for credentials, perform the following steps:

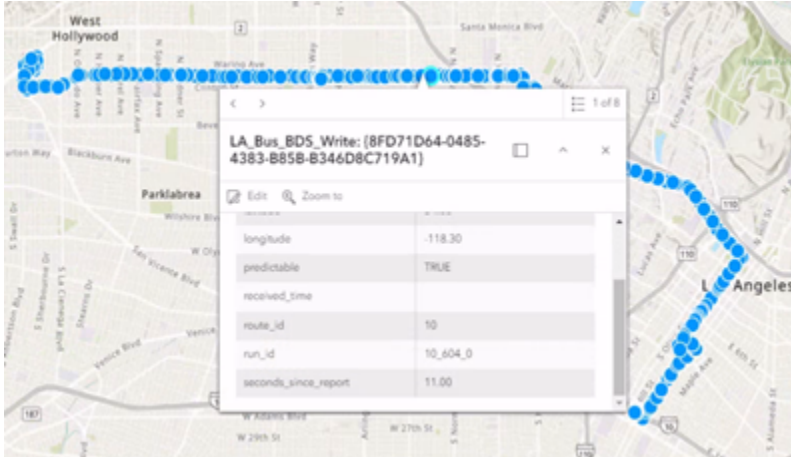
1. Sign in as **siteadmin** with the **Esri.4.GIS** password.
2. Repeat the process to delete the features using the By Where Clause with the SQL expression `route_id <> 10`.

The Manage Spatiotemporal Data Stores page will appear again after the data is deleted.

- o Restore the web browser tab in which Map Viewer is active.
- p In the Layers pane, for the LA Bus BDS Write feature layer, click the Options button  and choose Zoom To.

Hint: You may need to refresh the web browser for the deleted data to disappear.

- 9 In the map, click any of the points, and then in the pop-up window, verify that the Route_id is 10.



You successfully deleted all the stored events, except for the Route_id of 10, from the underlying spatiotemporal big data store data source.

- r Close the web browser without saving the map.

In this exercise, you investigated the spatiotemporal big data store installed on the STBDS VM. The data store contained about 100,000 bus locations that were published and managed from ArcGIS GeoEvent Manager as hosted services. Like all hosted services, the hosted services published from GeoEvent can be found on the hosting server, and layers can be visualized within the portal.

Distributed data and the spatiotemporal big data store

You can install the spatiotemporal big data store on multiple machines (nodes). All spatiotemporal big data store nodes need to be configured with the same GIS Server site, the hosting server. Adding additional machines provides the following benefits:

- Adds storage capacity, enabling you to store more data
- Allows for isolated coordination by specifying the role of each node
- Increases the data write throughput through the use of sharding
- Improves data availability through the replication factor

This method of distributed data also means that the failure of one machine will not cause a loss of data, thereby supporting high availability.



When adding machines to your spatiotemporal big data store, only an odd number of nodes (such as three or five) are supported.

Node role

The spatiotemporal big data store nodes can be assigned a role of data or coordinator.

- Data nodes are used to store data and are more memory intensive.
- Coordinator nodes manage cluster states, manage metadata, and are less memory intensive.
- Coordinator and data nodes have the responsibilities of both roles.



There can only be one, three, or five nodes with any type of coordinator role because this role must be deployed in odd numbers and assigned to a maximum of five nodes. These requirements do not apply to data nodes.

Sharding

A multi-node spatiotemporal big data store employs a type of database partitioning, called

Distributed data and the spatiotemporal big data store (continued)

sharding, to distribute and replicate data between nodes. Sharding is replicating the schema of the data and then dividing the data based on a shard key onto a separate spatiotemporal big data store instance to spread the load.

Replication

As an administrator, you can also determine how many times a dataset can be replicated, which provides redundancy that protects against data loss. Shards and replica shards are automatically distributed amongst nodes such that a primary shard and its replica shard will not be located on the same node.

Examples

In this example, a three-node configuration with no dedicated data nodes is shown.

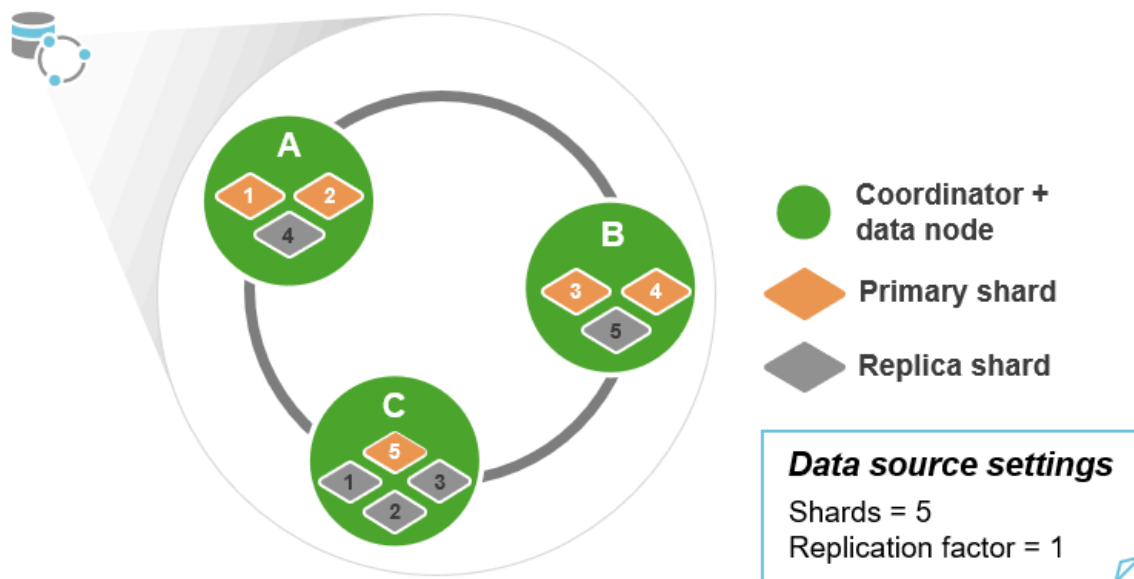


Figure 9.6. In this configuration, a single data source is split into five primary shards (orange icons) that are distributed across the three coordinator and data nodes. Each primary shard is replicated one time (gray icons) to another node.

Distributed data and the spatiotemporal big data store (continued)

In this example, a five-node configuration with dedicated nodes is shown.

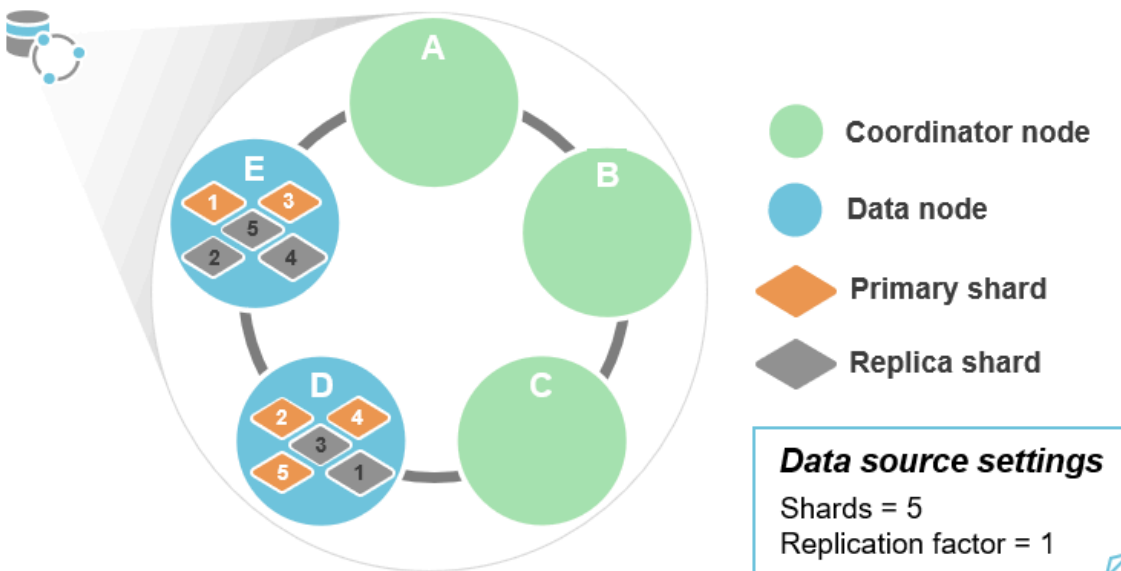


Figure 9.7. In this configuration, a single data source is split into five primary shards (orange icons) that are distributed across the two data nodes. Each primary shard is replicated one time (gray icons) to another data node.

Distributed analysis and the spatiotemporal big data store

ArcGIS GeoAnalytics Server uses a distributed computing framework to help users process high volumes of vector and tabular data quickly. Similar to raster analytics, GeoAnalytics tools leverage distributed data, enabling distributed analysis. With this framework, processes can be completed on a single machine or distributed across multiple machines.

GeoAnalytics Server can be installed on one machine for a single machine site or on three machines for a multiple-machine site. It is best practice to have the same number or a greater number of spatiotemporal big data store nodes than the number of machines in the GeoAnalytics Server site.

The special tools that GeoAnalytics Server exposes in the Enterprise portal Map Viewer Classic and in ArcGIS Pro allow users to perform big data analyses using the computing power of the GeoAnalytics Server machines.

When running GeoAnalytics Server tools, the analysis is completed on the GeoAnalytics Server. However, data that is not local to your GeoAnalytics Server will be moved to your GeoAnalytics Server before analysis begins, so it is recommended that you make data available to the GeoAnalytics Server through hosted feature layers or through big data file shares.

Distributed analysis and the spatiotemporal big data store (continued)

Distributed computing using GeoAnalytics Server and the spatiotemporal big data store

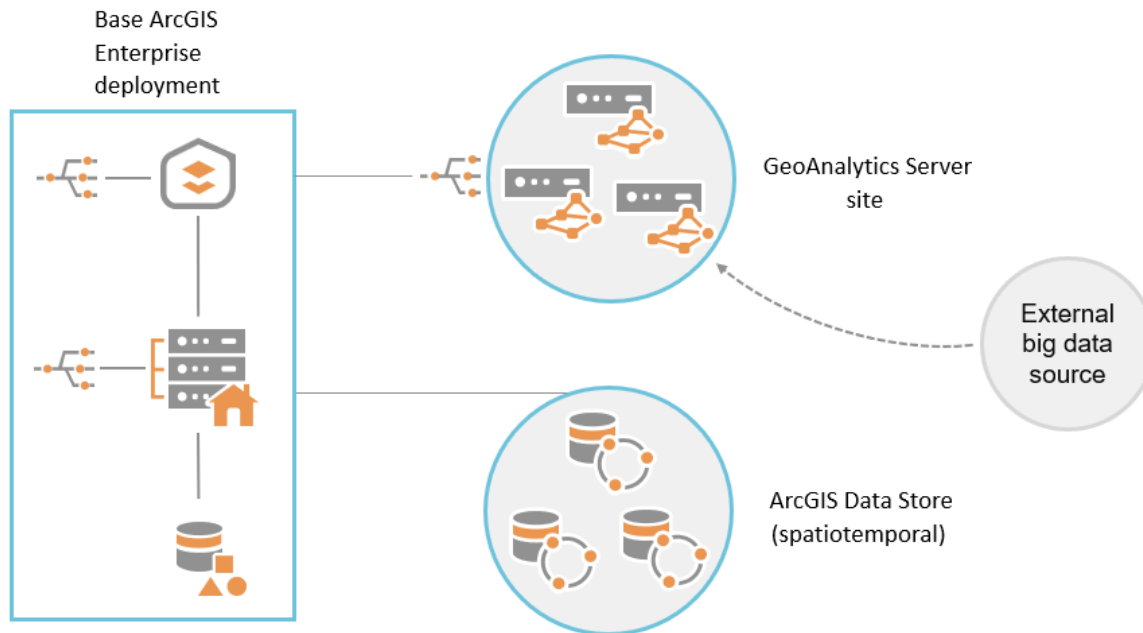


Figure 9.8. This example illustrates a three-machine configuration of GeoAnalytics Server. This configuration of the spatiotemporal big data store meets system requirements with three nodes; however, if necessary, additional nodes could be added.

GeoAnalytics Server requires the spatiotemporal big data store to store analysis results. Considering the size of potential output data and the resource-intensive nature of GeoAnalytics processing, you should carefully consider which Enterprise portal members to grant access to the tools. For members to be able to use GeoAnalytics capabilities, the following privileges must be enabled by an administrator:

- Create, update, and delete content
- Publish hosted feature layers
- GeoAnalytics Feature Analysis

Lesson review

1. The spatiotemporal big data store is a requirement to use ArcGIS GeoEvent Server.
 - a. True
 - b. False

 2. With which federated ArcGIS Server site is the spatiotemporal big data store registered?
 - a. GeoAnalytics Server
 - b. GeoEvent Server
 - c. Hosting server
 - d. Image Server

 3. ArcGIS GeoAnalytics Server can both read and write to a spatiotemporal big data store.
 - a. True
 - b. False

 4. Why might you want to deploy the spatiotemporal big data store across multiple nodes?
-
-

Answers to lesson 9 questions

How the spatiotemporal big data store relates to other ArcGIS Enterprise components

(page 9-8)

| Functions related to the spatiotemporal big data store | Match letter |
|--|--------------|
| The spatiotemporal big data store is registered with it. | B |
| It writes real-time observational data to the spatiotemporal big data store. | C |
| Additional spatiotemporal big data store nodes should be installed on it. | A |
| It can analyze data from the spatiotemporal big data store. | D |

Exercise 9: Explore data in the spatiotemporal big data store (page 9-9)

1. Which of these data stores do you believe is the spatiotemporal big data store?
_bigdata_bds is the spatiotemporal big data store.
2. Which port must be open for spatiotemporal BDS communication?
Database port 9320
3. What are the three types of server connections that can be made?
Connections can be made to ArcGIS Server, ArcGIS Enterprise, or ArcGIS Online.
4. Which URL is used to connect to ArcGIS Enterprise on EBASE?
The URL of your Enterprise portal (<https://ebase.ad.local/portal>)
5. If you wanted to view this hosted map service in ArcGIS Server Manager, which ArcGIS Server site would you use?
Because it is a hosted service, it will appear in ArcGIS Server Manager on the hosting server (ebase.ad.local/server), as indicated by the URL.

In the real world, things are constantly changing. Whether you are tracking moving assets, monitoring stationary sensors, or following social media feeds, there are special considerations for incorporating these real-time changes into your ArcGIS Enterprise deployment. ArcGIS GeoEvent Server is designed to work with the unique characteristics of real-time data streams to ingest and process data quickly and efficiently. The end result is that real-time data can be integrated into workflows, analyses, and applications more than ever before.

Topics covered

Understanding real-time data

Filtering and processing real-time data in ArcGIS GeoEvent Server

Building a GeoEvent Service

Real-time data

Organizations gain situational awareness when they monitor relevant real-time data streams. Real-time data originates from various sources, but observations of events happening in real time can be grouped into three categories.

Categories of real-time observations

Stationary sensors



Things that move



Things that "just happen"



What are some examples of real-time observational data?

Working with real-time data

ArcGIS GeoEvent Server enables real-time event-based data streams to be integrated as data sources in your enterprise GIS. GeoEvent Server has five key capabilities to work with real-time observational data.



Ingestion: The first step to working with real-time data in your ArcGIS Enterprise deployment is to get the data into your system. GeoEvent Server includes configurable inputs to connect to and ingest various types of data.



Real-time analytics: GeoEvent Server can filter high-velocity data as it passes through the system. Filters can be applied to attributes of the data as well as spatial relationships (for example, a geofence). Processing can also occur in real time. For example, fields can be calculated and buffers created as GeoEvents pass through the system.



Data archiving: After data has been filtered and processed, the data can be stored as feature records in a traditional relational geodatabase or ArcGIS Data Store. The spatiotemporal big data store, for example, can persist huge amounts of real-time data as it streams into the system.



Visualization: Enhanced map and feature services aid in visualizing feature records. For example, features can be aggregated on demand. Stream services allow applications to connect directly to GeoEvent outputs through a WebSocket.



Notification and control: Configurable outputs provide the capability to format an email or SMS text message. GeoEvent Server can send data to external IoT interfaces to control devices.

What is a GeoEvent?

When ArcGIS GeoEvent Server connects to a data feed, it receives reports broadcasted by the feed. Regardless of whether the report contains information about moving vehicles, stationary sensors, or new traffic accidents, these reports are made up of individual records that describe each event or observation. A *GeoEvent* is one of these event records with location information attached.

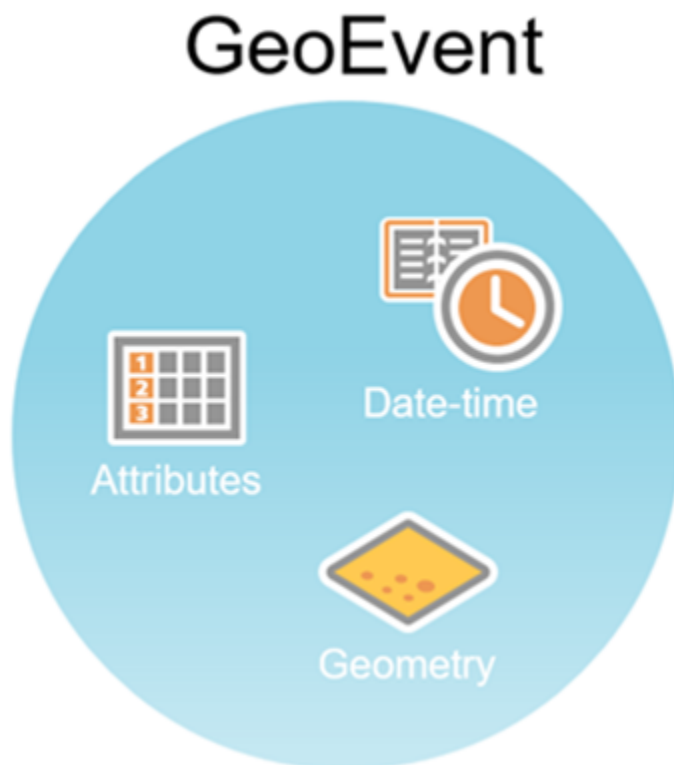


Figure 10.1. A GeoEvent includes information about an incident or observation, as well as when and where it occurred.

A key aspect of GeoEvent Server is the ability to work with data as it streams through the system. In the human body, blood cells are enriched with oxygen, and waste is removed as the cells travel through capillaries. Conceptually, a GeoEvent Service is similar in that the service acts on GeoEvents as the GeoEvents pass through the system in real time. GeoEvents that are ingested can be manipulated in various ways to add value. The following list provides some examples of actions that can be applied to real-time data:

What is a GeoEvent? (continued)

- Filter GeoEvents using spatial or attribute conditions to focus on the most interesting event data.
- Enrich incoming events with data from a secondary feature service or system file.
- Calculate new fields in real time.
- Create new geometries based on incoming events such as buffering features or creating points from polygons.




Figure 10.2. Because GeoEvent records contain geometry and attribute information, the GeoEvent Server can further filter and process real-time data as it is ingested.

Explore configuring real-time analytics

Using ArcGIS GeoEvent Server, you can build GeoEvent Services that define how to connect to a real-time data stream, how to manipulate the data, and what to do with the result.

You will use ArcGIS Enterprise help documentation to explore several key aspects of GeoEvent Services.

Instructions

- a** Open a web browser and go to the [ArcGIS GeoEvent Server: Introduction To GeoEvent Services](https://enterprise.arcgis.com/en/geoevent/11.3/geoevent-services/overview-of-geoevent-services.htm) web page (https://enterprise.arcgis.com/en/geoevent/11.3/geoevent-services/overview-of-geoevent-services.htm).
- b** Review this page to answer the first question in your workbook.
- c** Click the [Processor](https://enterprise.arcgis.com/en/geoevent/11.3/analyze/processors.htm) link (https://enterprise.arcgis.com/en/geoevent/11.3/analyze/processors.htm), and then answer the second question in your workbook.
- d** After you are finished, close the web browser.

1. Which components are combined to form a GeoEvent Service?

2. What is an example of a processor that is applicable to your work?

Building a GeoEvent Service in GeoEvent Manager

GeoEvent services are built in the GeoEvent Manager service designer—an interface similar to ModelBuilder—within GeoEvent Manager. These services are made up of four components: inputs, filters, processors, and outputs.

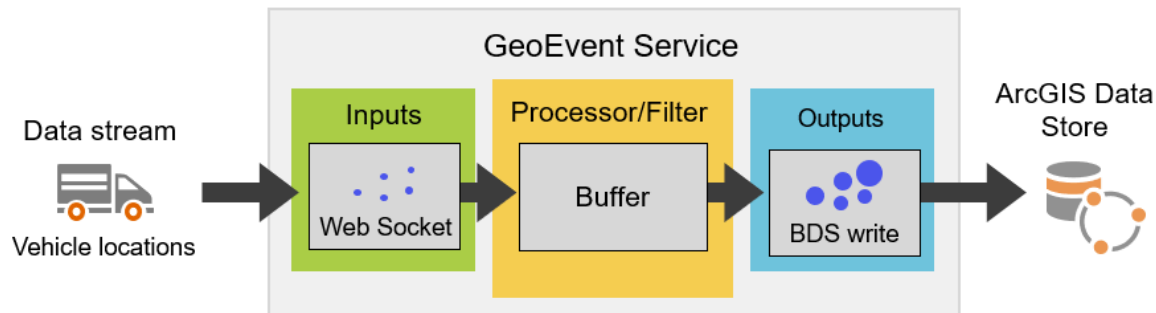


Figure 10.3. The components of a GeoEvent Service represent the flow of GeoEvents through ArcGIS GeoEvent Server—where the data comes from, what happens to it, and where it goes.

The process for building a GeoEvent Service includes the following steps:

1. Add an input and output to the canvas. (A GeoEvent Service requires one input and at least one output, although it can include multiple outputs.)
2. Optionally, add any filters or processors that you need to manipulate the data in real time.
3. Connect the input to all outputs, routing through any filters or processors that you have included.
4. Publish the service to begin the real-time processing of the event data.

Exercise 10

Configure a GeoEvent Service to write real-time data to a

In the previous exercise, you were introduced to GeoEvent Manager. You will now use the service designer in GeoEvent Manager to publish a new GeoEvent Service.

Your goal is to begin storing real-time location data of buses in Los Angeles, California, in a spatiotemporal big data store. The end result will be a web layer that your Enterprise portal members can use to visualize the current state of the data while using features such as dynamic aggregation. You will use GeoEvent Simulator to simulate a real-time data feed as you build the GeoEvent Service, which can be updated in the future to connect to a real-time feed after you receive feedback from your stakeholders.

In this exercise, you will perform the following tasks:

- Federate a GeoEvent Server site.
- Explore a configured input connector and GeoEvent Definitions.
- Create an output connector.
- Publish a GeoEvent Service.

Step 1: Federate a GeoEvent Server site

To begin, you will federate the GeoEvent Server site with your ArcGIS Enterprise deployment.

To leverage GeoEvent Server with ArcGIS Enterprise, federation is not required. In this exercise, you will federate to incorporate ArcGIS Enterprise portal's security model to access GeoEvent Manager with the portaladmin credentials.

- a Using the Remote Desktop Connection Manager, in the server tree panel, double-click EBASE to connect to the EBASE VM.
- b Sign in to the adAdmin account (password: **Esri.4.GIS**), if necessary.



The desktop for the EBASE VM is blue.

- c In File Explorer, browse to `\\EADM\EsriTraining\EADM\GeoEvent`.

This folder includes a text file that contains the URL that you will copy to ensure accuracy as you federate the server.

- d Open the FederateGeoEventServer.txt file and leave it open for the remainder of the step.
- e Open a web browser and click the ArcGIS Enterprise bookmark, if necessary.
- f Sign in to the Enterprise portal as **portaladmin** (password: **Esri.4.GIS**).
- g Click Organization, and then click the Settings tab to access the ArcGIS Enterprise settings page.
- h On the left side, click the Servers tab.
- i Under Federated Server Sites, click Add Server Site, and then specify the following information:

- Services URL: **https://geo.ad.local/geo** (copy and paste from the text file)
- Administration URL: **https://geo.ad.local:6443/arcgis** (copy and paste from the text file)
- Username: **siteadmin**
- Password: **Esri.4.GIS**

j Click Next.



Federating the server may take several minutes.

k After you progress to the Configure Server Role page, scroll through the available server roles.

There is not a GeoEvent Server role because GeoEvent Server has no special capabilities gained through federation. However, federating your GeoEvent Server site means that it will use the same security model as your Enterprise portal.

l Click Done.

In a moment, the federated server will appear at the bottom of the list of federated server sites.

| | | |
|-----------------------|---|---|
| ▼ geo.ad.local:6443 ✎ | | ⋮ |
| Status | ✔ All systems operational | |
| Server role | Configure server role | |
| Service URL | https://geo.ad.local/geo | |
| Administration URL | https://geo.ad.local:6443/arcgis | |

Your GeoEvent Server site is now federated with your Enterprise portal.

Step 2: Investigate an existing input connector

Next, you will investigate an existing input connector using GeoEvent Manager. This input connector has already been configured to connect to a text file that contains bus locations in Los Angeles, California. You will use this input connector to create a GeoEvent Service in a subsequent step.

- a** Using the Remote Desktop Connection Manager, connect to the GEO VM (light brown desktop).
- b** Sign in to the adAdmin account (password: **Esri.4.GIS**), if necessary.

- c Open a web browser and click the GEO-GeoEvent Manager bookmark.



If the GeoEvent Manager web page does not appear, stop the following services on the GEO machine in the following order: ArcGIS GeoEvent Server service, ArcGIS GeoEvent Gateway service, and ArcGIS Server service. Then, restart the ArcGIS GeoEvent Server service. (Restarting this service will trigger the other two services to restart.)

- d Sign in as **portaladmin** (password: **Esri.4.GIS**).

You are able to sign in using your Enterprise portal administrator credentials because you federated the GeoEvent Server site.

- e Click the Manager tab, if necessary.

You are now viewing the Manager page. The Manager page displays the status and activity of input connectors, GeoEvent Services, and output connectors in the GeoEvent site.

- f Under the Inputs section, click the link for Tcp-text-in.

You are now viewing the properties of this input connector. This input was created using one of the standard input connectors that comes with GeoEvent Server. It is configured to receive text from a TCP socket. If you wanted to receive data from a different source, such as JSON data from an external server's WebSocket, you would use a different input connector. GeoEvent Server comes with input connectors to connect to various data sources.

- g In the input connector dialog box, expand the Advanced section, if necessary.

1. Which server port will be used to accept TCP client connections?


You will use this port in the next step to send data to this input connector.

- h Close the input connector dialog box.

Step 3: Configure the GeoEvent Simulator

Next, you want to test the tcp-text-in input connector using GeoEvent Simulator before continuing to build the GeoEvent Service.

- a To the right of Tcp-text-in, locate the Start button , as indicated in the following graphic.




- b Click the Start button  to start the input connector.

The input connector has started, but no data is being sent to it.


- c From the Windows Taskbar, start GeoEvent Simulator by clicking the GeoEvent Simulator icon, as indicated in the following graphic.



- d Resize your web browser window and position it next to the GeoEvent Simulator window so that you can see both applications.
- e In GeoEvent Simulator, for Server, verify that the hostname is set to localhost and the TCP port is set to 5565.



Because GeoEvent Server is running on this machine, you will use these default settings. You previously noted that your input connector is listening on port 5565 to receive data. Before connecting to the tcp-text-in connector, you will load your event data, which is stored as a CSV file.

- f Click the Load File button .
- g In the Load From File dialog box, for File, browse to `\\EADM\EsriTraining\EADM\GeoEvent` and double-click the LA-Bus_Data.csv file.

The Load From File dialog box will display a preview of the data. Here, you will confirm that GeoEvent Simulator will correctly parse the data into fields. You will identify which field contains the time field in the data. Also, because the first row of the CSV file identifies field names, you will tell GeoEvent Simulator to skip the first row of the file when you choose to send data to the input connector.

- h In the Load From File dialog box, set the Time Field # parameter to **0**.

i Check the box for Skip The First 1 Lines.

File Properties

File: \\EADM\EsrTraining\EADM\GeoEvent\LA-Bus_Data.csv

Encoding: Auto Detect

Event Separator: \n

Field Separator: .

Time Field # 0 11/27/2021 02:25:30 PM

Auto Detect Time Field Epoch (up to 10-digits) Seconds Value

Skip the First 1 Lines

Preview Events

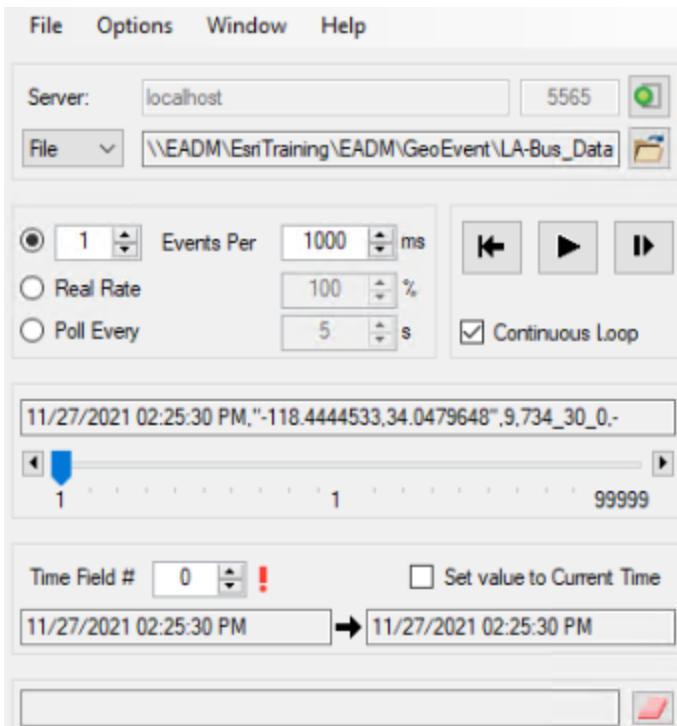
Show Event Separator Fields View


| Field 0 | Field 1 | Field 2 | Field 3 |
|------------------------|---------------------------|---------|---------|
| 11/27/2021 02:25:30 PM | "-118.4444533,34.0479648" | 9 | 734 |
| 11/27/2021 02:25:30 PM | "-118.4660585,34.1627555" | 37 | 734 |
| 11/27/2021 02:25:30 PM | "-118.4662475,34.155251" | 93 | 734 |
| 11/27/2021 02:25:30 PM | "-118.449257,34.274487" | 93 | 734 |
| 11/27/2021 02:25:30 PM | "-118.436241,34.034908" | 178 | 734 |
| 11/27/2021 02:25:30 PM | "-118.4504122,34.2928357" | 9 | 734 |
| 11/27/2021 02:25:30 PM | "-118.453453,34.077038" | 207 | 734 |
| 11/27/2021 02:25:30 PM | "-118.4676972,34.2386013" | 9 | 734 |

j In the Load From File dialog box, click Load.

Now you are ready to connect to the input connector and send data.

k Click the Click To Connect button .




- l** In GeoEvent Simulator, click the Step button  five times to send some records to the input connector.
- m** In your web browser tab that shows GeoEvent Manager, look at the summary information for the Tcp-text-in connector.

The Count field should be updated to indicate that several event records have been received.



The Count field may take a couple minutes to update; there is no need to refresh the web browser.

- n** Click the Step button  several more times, if you would like, and observe the activity in the Tcp-text-input connector.

You confirmed that your input connector is configured.


Step 4: Investigate GeoEvent Definitions

Before continuing to build your new GeoEvent Service, you will investigate GeoEvent Definitions to gain an understanding of how schemas are handled in GeoEvent Server.

- a** In GeoEvent Manager, click the Site tab.

GeoEvent Definitions are viewed and managed from the Site tab.

- b** Confirm that the GeoEvent Definitions tab is active.
- c** In the list of GeoEvent Definitions, locate the row for LA_Bus_CSV_Input.



This GeoEvent Definition is the one that is specified in your input connector.

- d** For the LA_Bus_CSV_Input GeoEvent Definition, click the Edit button  .

The fields that are listed in this GeoEvent Definition are consistent with the data in the CSV file that you loaded previously, which is necessary for the input connector to construct GeoEvents. You can also see which fields have associated tags, such as TRACK_ID, GEOMETRY, and TIME_START. These are examples of built-in tags that identify fields with critical information.

- e** Click Cancel to return to the list of GeoEvent Definitions.
- f** In the list of GeoEvent Definitions, locate the row for LA_Bus_BDS_Write.

To send the event data to the spatiotemporal big data store, you must also have a GeoEvent Definition that is consistent with the schema of the output data source. You will use the LA_Bus_BDS_Write GeoEvent Definition for this purpose in a subsequent step.

Step 5: Create a new output connector

Now you are ready to create an output connector that will write GeoEvent data to the spatiotemporal big data store.

- a** Click the Manager tab.
- b** Under the Outputs section, click Add Output.
- c** Review the list of output connectors and read several descriptions.

Various output connectors are available to send events to a single specific output. Every GeoEvent Service must include at least one output connector.

- d** Under Spatiotemporal Big Data Stores, click the link for the Add A Feature To A Spatiotemporal Big Data Store connector.

- e Configure the output connector using the following parameters:
- For Name, type **NEW-bds-out**.
 - For Registered Server Connection, from the drop-down list, choose ArcGIS Enterprise On EBASE.

Creating Output - Add a Feature to a Spatiotemporal Big Data Store Save Cancel Help

Name*:

Registered server connection*: Register ArcGIS Server

Data Source Name*: Create Data Source



Do not click Save after specifying these parameters.

This output connector also requires a place to send data. You will create a new data source in the spatiotemporal big data store to archive the data. You can create a data source on the Site tab, but you can also create a data source here for convenience.

- f Click Create Data Source.
- g In the Create Data Source window, for Name, type **NEW_LA_Bus_BDS_Write**.

Create Data Source

Name:

GeoEvent Definition:

You determined that there is no need to create a feature service associated with this data source, so you will only create a map service.

- h Expand the Service section, if necessary, and then for Create, uncheck the box for Feature Service.

▼ Service

Create: Map Service Feature Service

- i Expand the Advanced section and change the Replication Factor to **0** (zero).

▼ Advanced

Replication Factor:

Because you are using a single-node spatiotemporal big data store, there can be no replication to other nodes.

- j** Expand the Aggregation Rendering section and set the Feature Threshold to **500**.

▼ Aggregation Rendering

Feature Threshold:

- k** Scroll down and click Publish.

Name*:

Registered server connection*: [Register ArcGIS Server](#)

Data Source Name*: [Create Data Source](#)

▶ Advanced

After the new data source is published, you are returned to the Add Output dialog box.

- l** In the Add Output dialog box, click Save to create the output connector.

▼ Outputs ▶ ■

[Add Output](#)

[LA_Bus-bds-out](#)

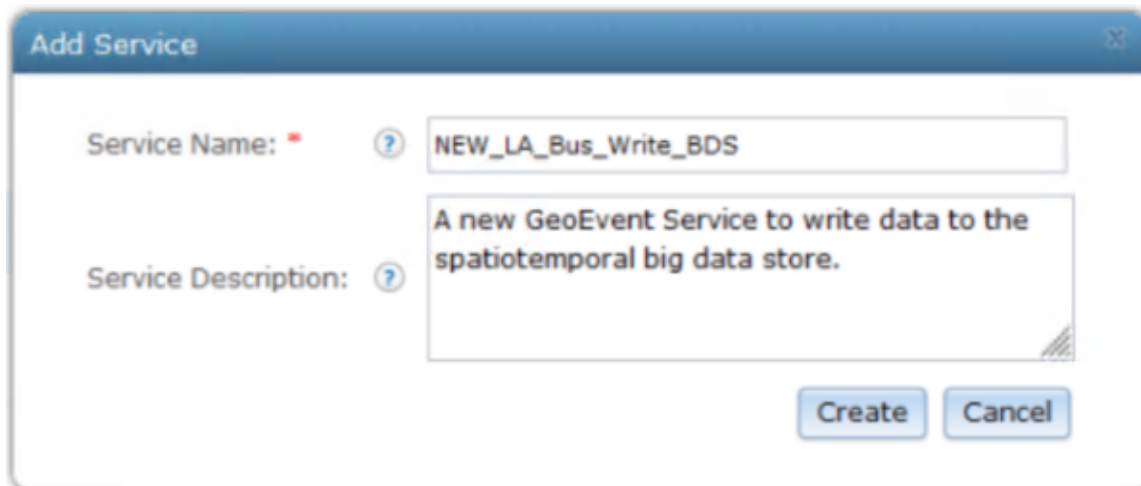
[NEW-bds-out](#)

Your output connector now appears in the Outputs section.

Step 6: Create a GeoEvent Service

Now that you have created an output connector, you can link the input connector that you examined earlier with your new output connector to begin writing data to the spatiotemporal big data store from the data stream.


- a From the Manager tab, under the GeoEvent Services section, click Add Service.
- b In the Add Service window, specify the following parameters:
 - Service Name: **NEW_LA_Bus_Write_BDS**
 - Service Description: **A new GeoEvent Service to write data to the spatiotemporal big data store.**



- c Click Create.

The service designer opens with a blank canvas. The service designer interface is similar to ModelBuilder in ArcGIS Pro; however, instead of working with geoprocessing elements, you work with GeoEvent Service elements (inputs, outputs, filters, and processors).

- d In the pane on the left side of the screen, expand Inputs, if necessary, and then drag Tcp-text-in onto the canvas.



- e Under Outputs, drag NEW-bds-out onto the canvas.
- f Under New Elements, double-click Processor.

g In the Processor Properties dialog box, specify the following parameters:

- For Name, type **LA_Bus_Field_Map**.
- For Processor, from the drop-down list, choose Field Mapper.
- For Source GeoEvent Definition, choose LA_Bus_CSV_Input.
- For Target GeoEvent Definition, choose LA_Bus_BDS_Write.

To complete the Field Mapper processor, you must match each source field to its corresponding field in the target GeoEvent Definition. Most of the fields have the same names and will be automatically matched, but you will complete the process for the Received_time target field.

h Under Source Fields, at the bottom of the list, click the down arrow that corresponds to Received_time and choose Event_time.

i Verify that the Source Fields and Target Fields match the following graphic.

| Source Fields | Target Fields |
|----------------------|------------------------------------|
| geometry | geometry <i>Geometry</i> |
| seconds_since_report | seconds_since_report <i>Double</i> |
| run_id | run_id <i>String</i> |
| longitude | longitude <i>Double</i> |
| heading | heading <i>Double</i> |
| route_id | route_id <i>String</i> |
| predictable | predictable <i>String</i> |
| latitude | latitude <i>Double</i> |
| id | id <i>String</i> |
| event_time | received_time <i>Double</i> |

In this case, the fields are very similar; however, you would use the same process to map fields between very different source and target schemas. For example, if your source data uses a nested field structure, such as JSON data, you would use the Field Mapper processor to map

the fields to a schema that uses a flat structure, which is necessary for storing data in the spatiotemporal big data store.

- j Click OK.
- k Reposition the elements so that the Field Mapper processor is positioned between the input and the output.



- l Point to the right edge of the Tcp-text-in input connector until the arrow icon appears, and then drag the arrow icon to the Field Mapper processor.
- m Repeat the process to connect the Field Mapper processor to the output.



- n Right-click the NEW-bds-out connector and choose Start Output.

Green start icons will indicate that both your input and output connectors are started. You are now ready to publish this GeoEvent Service, which links the input and output connectors.

- o Click Publish.
- p Notice the NEW_LA_Bus_Write_BDS service status, as indicated in the following graphic.

NEW_LA_Bus_Write_BDS

A new GeoEvent Service to write data to the spatiotemporal big data store.

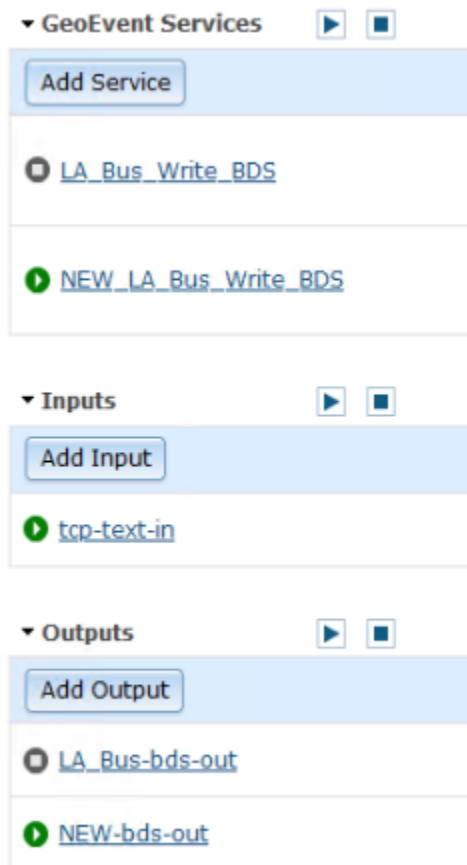
| Status | In/Out | Count | |
|----------------|--------|-------|--|
| STARTED | In | 0 | |
| | Out | 0 | |

When you see that the `NEW_LA_Bus_Write_BDS` services status indicates that it has started, you know that the GeoEvent Service has been created.

Step 7: Observe data as it flows through a GeoEvent Service


As a final step, you will use the Manager page and the map image layer to confirm that data is being written to the spatiotemporal big data store.

- a Click the Manager tab.
- b Confirm that the following elements are started:
 - `NEW_LA_Bus_Write_BDS` (under GeoEvent Services)
 - `Tcp-text-in` (under Inputs)
 - `NEW-bds-out` (under Outputs)





Now you are ready to test the GeoEvent Service by sending data to the input connector. You will

use GeoEvent Simulator like you did previously to control when events are sent.

- c As you did previously, position your web browser window and the GeoEvent Simulator window side by side.
- d In GeoEvent Simulator, click the Step button  several times to send event data to the input connector.



If the Count field increases for the input connector but not the GeoEvent Service and output connector, perform the following steps to restart GeoEvent Server:

1. On the GEO VM, click the Windows Start icon and search for **services**, and then open the Services app.
2. Right-click and choose Stop for each of the following services (in the following order): ArcGIS GeoEvent Server, ArcGIS GeoEvent Gateway, and ArcGIS Server service. (This action stops GeoEvent Server.)
3. When the services are stopped, right-click the ArcGIS GeoEvent Server service and choose Start. (Restarting this service will trigger the other two services to restart.)
4. Close the Services app and wait a couple minutes before refreshing the web browser.
5. When GeoEvent Manager is restarted, notice that the Counts will return to 0, which is expected.
6. Open GeoEvent Simulator, click the Click To Connect button  to restore the connection to localhost, and then click the Step button  several times to send event data to the input connector.


2. Are events streaming through the system? How can you tell?

You will now set up a map that will allow you to visualize the data that is stored in the spatiotemporal big data store.


- e Open a new web browser tab and click the ArcGIS Enterprise bookmark.
- f Sign in as **portaladmin** (password: **Esri.4.GIS**), if necessary.
- g Click Content.
- h Locate the NEW_LA_Bus_BDS_Write map image layer and add it to a new map.

Hint: Click the More Options button **⋮** and choose Open In Map Viewer.

- i From the Properties pane on the right, scroll down to the Refresh Interval section.
- j Enable Automatically Refresh Layer.
- k Set the refresh interval to **0.1** minutes.
- l Click in the white space of the pane to apply the new setting.




- m In the map, zoom to the Los Angeles, California, area.






The number of points that you see on your map will vary, depending on how many events that you sent through using GeoEvent Simulator.

- n In GeoEvent Simulator, change the rate of events to **100** Events Per 1000 Ms.



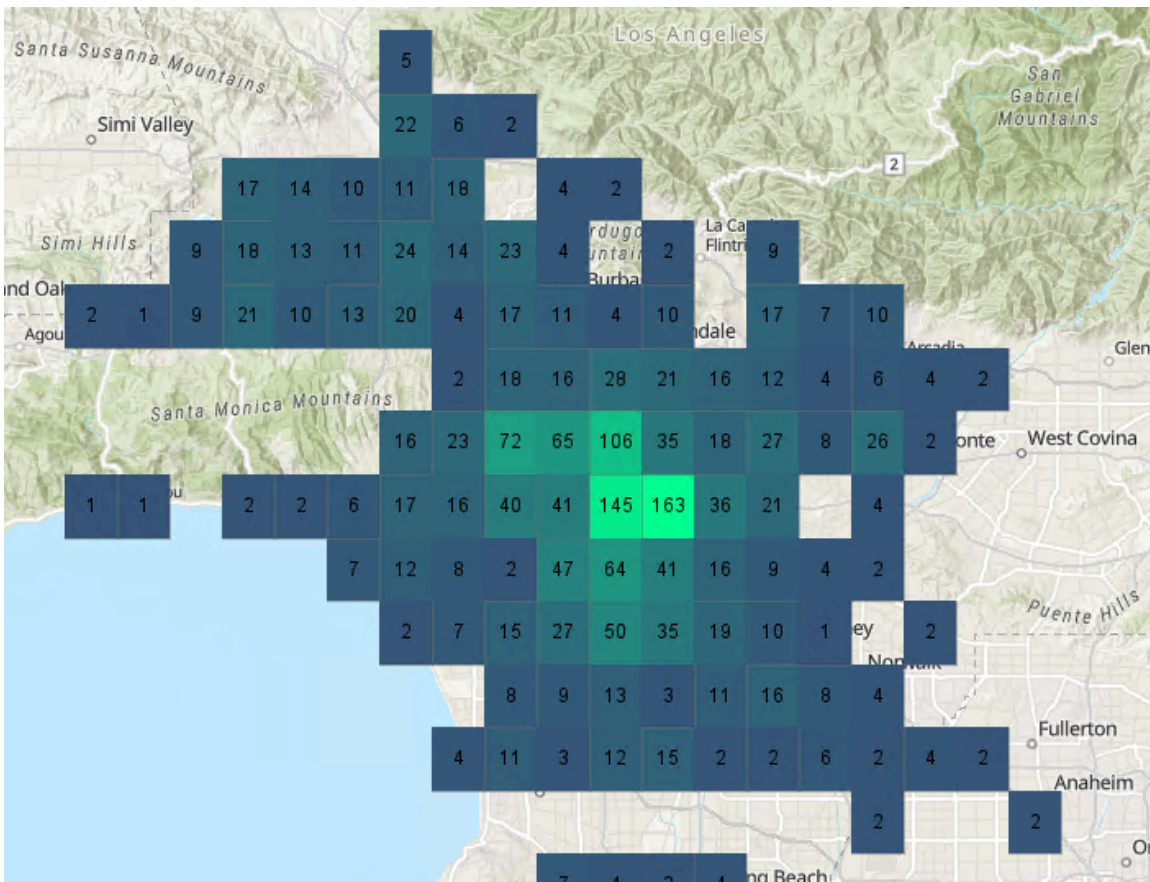
- o In GeoEvent Simulator, click the Play button .

- p Observe your map of Los Angeles, California.





If the map does not update within six seconds, refresh the browser.

- q Observe the map until the features begin to be aggregated.



When you published the map service, you set an aggregation threshold of 500. This threshold

setting means that when more than 500 features must be rendered in the display, they will be aggregated according to the aggregation style that is specified for the service (for example, geohash, square, and hexagon).

- r** In GeoEvent Simulator, click the Pause button  to stop the stream of data.
- s** Restore the ArcGIS GeoEvent Manager web browser tab, and then for the NEW_LA_Bus_Write_BDS GeoEvent Service, click the Stop button .

| | | |
|--|-----|-------|
|  NEW_LA_Bus_Write_BDS | In | 2,792 |
| | Out | 2,792 |

- t** Notice the number in the Count field for NEW-bds-out that indicates the number of data points that were ingested into the spatiotemporal big data store before stopping the stream of data.



If you do not pause the GeoEvent Simulator, you may negatively affect the hardware resources available to your system in subsequent exercises.

- u** Close the web browser and exit GeoEvent Simulator.

Although federation is not required to use GeoEvent Server with ArcGIS Enterprise, in this exercise, you federated GeoEvent Server to incorporate ArcGIS Enterprise portal's security model. Next, in GeoEvent Manager, you investigated an existing input connector that connected to a text file containing bus locations in Los Angeles, California. After creating a new output connection, you used the existing input and new output connector to publish a GeoEvent Service. Using GeoEvent Simulator, you observed the ingestion, processing, and output of GeoEvent records in GeoEvent Manager and visualized the layers in portal.

Lesson review

1. A GeoEvent Service can have more than one output.
 - a. True
 - b. False

2. Sending an email notification when a vehicle leaves a certain geographic area is a task performed by which component of a GeoEvent Service?
 - a. Processor
 - b. Output connector
 - c. GeoEvent Definition
 - d. Filter

3. Which element of a GeoEvent Service is used to create buffers around GeoEvents after they are ingested into GeoEvent Server?
 - a. GeoEvent Simulator
 - b. Filter
 - c. GeoEvent Definition
 - d. Processor

Answers to lesson 10 questions

Real-time data (page 10-2)

What are some examples of real-time observational data?

Possible responses may include the following:

- **Stationary sensors:** stream gauges, weather stations, air quality sensors, temperature readings, utility status readings
- **Things that move:** automobiles, aircraft, ships, people, satellites
- **Things that "just happen":** crimes, lightning, accidents, tweets

Explore configuring real-time analytics (page 10-6)

1. Which components are combined to form a GeoEvent Service?

An input connector and output connector are combined to form a GeoEvent Service.

Optionally, filters and processors can be applied between inputs and outputs.

2. What is an example of a processor that is applicable to your work?

Answers will vary, but processors can perform tasks such as calculating fields, creating buffers around features, and mapping fields. Custom processors can be developed using the GeoEvent Server SDK.

Exercise 10: Configure a GeoEvent Service to write real-time data to a spatiotemporal big data store (page 10-8)

1. Which server port will be used to accept TCP client connections?

5565

2. Are events streaming through the system? How can you tell?

Events are streaming through the system, as indicated by the increasing numbers in the Count column for the GeoEvent Service and input/output connectors.

In previous lessons, you explored workflows and best practices of sharing and managing layers and services within ArcGIS Enterprise. Now you may begin to consider a broader architectural pattern that allows collaboration to take place between multiple distributed portals. With distributed collaboration, you can extend the reach of your GIS content by seamlessly sharing maps, apps, layers, and more with other ArcGIS organizations.

Think of distributed collaboration as a way that two or more organizations, departments, or business units can precisely share the information that they must share. It is common for such entities to have sensitive information that they must keep secure while allowing certain information to be readily shared. Distributed collaboration supports this objective without any coding or adding unfamiliar workflows for users.

Topics covered

- Patterns of distributed collaboration
- Sharing content in a collaboration
- Creating and managing a distributed collaboration

Sharing content between portals

Sharing geographic data and content has evolved over time. Before the proliferation of the internet, options for sharing GIS data were limited. As technology evolved, new ways to share geographic information were enabled through the use of web servers. In turn, the Web GIS paradigm was developed; members of an organization could now easily create and share content, as well as discover and use the content shared by other members of their organization. Additionally, there may be a potential need to collaborate with other organizations.

Think about your own organization's sharing needs to answer the following question.

What situations in the work that you do could benefit from sharing content between portals?

Distributed collaboration

When organizations that use separate ArcGIS Enterprise deployments have common goals or initiatives, distributed collaboration provides a way to connect multiple Enterprise portals together to support data access and sharing.

Each portal retains its own security model while using the same group sharing model used to share content internally. Content is easily shared by adding the content to a group that is associated with the target workspace. Updates to data are kept in sync automatically.

There are various patterns within distributed collaboration that involve ArcGIS Enterprise and ArcGIS Online.

Examples of collaboration patterns supporting business needs

The following table presents several examples of how a business need might be supported by creating a collaboration between one or more portals.





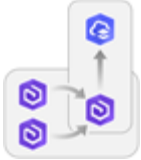
| Business need | Diagram | Collaboration pattern |
|--|---|--|
| An organization manages apps and data that need to be accessed by contractors and, eventually, the public. | ArcGIS Enterprise sharing to ArcGIS Online  | A collaboration is created to share data between the organization's ArcGIS Enterprise deployment and ArcGIS Online. |
| Data is collected in ArcGIS Online using ArcGIS Field Maps and public surveys, but members of an ArcGIS Enterprise portal must be able to access the data. | ArcGIS Online sharing to ArcGIS Enterprise  | Instead of creating ArcGIS Online member accounts for ArcGIS Enterprise users, a collaboration is created to share data from ArcGIS Online to ArcGIS Enterprise. |

Table continued on next page

Distributed collaboration (continued)

| Business need | Diagram | Collaboration pattern |
|---|--|--|
| <p>A project involving multiple agencies would benefit from a centralized, up-to-date view of data.</p> | <p>Multiple ArcGIS Enterprise portals sharing to ArcGIS Online</p>  | <p>A collaboration is created in which multiple ArcGIS Enterprise portals can share data to a centralized portal.</p> |
| <p>An organization has several ArcGIS Enterprise deployments and needs to push data and apps from a central office to multiple satellite offices.</p> | <p>An ArcGIS Enterprise portal sharing to multiple ArcGIS Enterprise portals</p>  | <p>A collaboration is created in which a single ArcGIS Enterprise portal can share data to multiple Enterprise portals.</p> |
| <p>In a temporary project, authoritative data must be shared with the lead agency, which then uses ArcGIS Online to share curated products with the public.</p> | <p>Two ArcGIS Enterprise portals sharing to an ArcGIS Enterprise portal that shares to ArcGIS Online</p>  | <p>Two collaborations are created. The first collaboration is used to share authoritative data from several agencies to the lead agency; the second collaboration is created so that the lead agency can use ArcGIS Online to share content with the public.</p> |



Planning for a collaboration

When planning for a collaboration between portals, it is important to understand the considerations that are associated with each element in a collaboration.


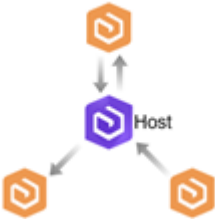


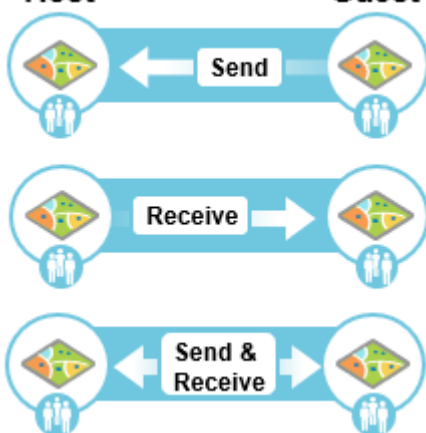
| Element | Considerations |
|---|---|
| <p>Host</p>  | <ul style="list-style-type: none"> • Creates the collaboration • Generates guest invitations and accepts their responses • Creates workspaces • ArcGIS Online can only serve as a host • ArcGIS Enterprise can be a host or a guest |
| <p>Guest</p>  | <ul style="list-style-type: none"> • Accepts invitation and generates response • Can contribute as a content originator • Portal that initiates content synchronization • ArcGIS Enterprise can be a guest or a host |
| <p>Workspace</p>  | <ul style="list-style-type: none"> • Connects collaboration groups to shared content • Created by the host and joined by guests • The host and each guest associate one group with a workspace • Defines how feature layers are shared and the type of editing that is supported (one way or two way) |
| <p>Collaboration group</p>  | <ul style="list-style-type: none"> • User interface for sharing collaboration content • Uses standard portal group sharing model • Badges identify collaboration-curated items • New copied collaboration content is owned by the collaboration group owner |

Table continued on next page

Planning for a collaboration (continued)

| Element | Considerations |
|---|--|
| <p>Access mode</p> <p>Host Guest</p>  <p>The diagram illustrates three access modes between a Host and a Guest portal. Each mode is represented by a horizontal bar with a Host icon on the left and a Guest icon on the right. 1. 'Send' mode: A white arrow points from the Guest icon to the Host icon, with the word 'Send' in a white box in the center. 2. 'Receive' mode: A white arrow points from the Host icon to the Guest icon, with the word 'Receive' in a white box in the center. 3. 'Send & Receive' mode: Two white arrows point in opposite directions (Host to Guest and Guest to Host), with the words 'Send & Receive' in a white box in the center.</p> | <ul style="list-style-type: none"> • Defined by the host for each workspace • Can be send only, receive only, or send and receive • Data can be configured to either only allow editing by the portal that sent the data or to allow two-way editing by both portals; to enable two-way editing, the guest's access mode must be set to send and receive • Access mode can vary for different guests |

Steps to create a distributed collaboration

After planning and understanding the components of a collaboration, there are a series of steps that should be completed to successfully create the collaboration.

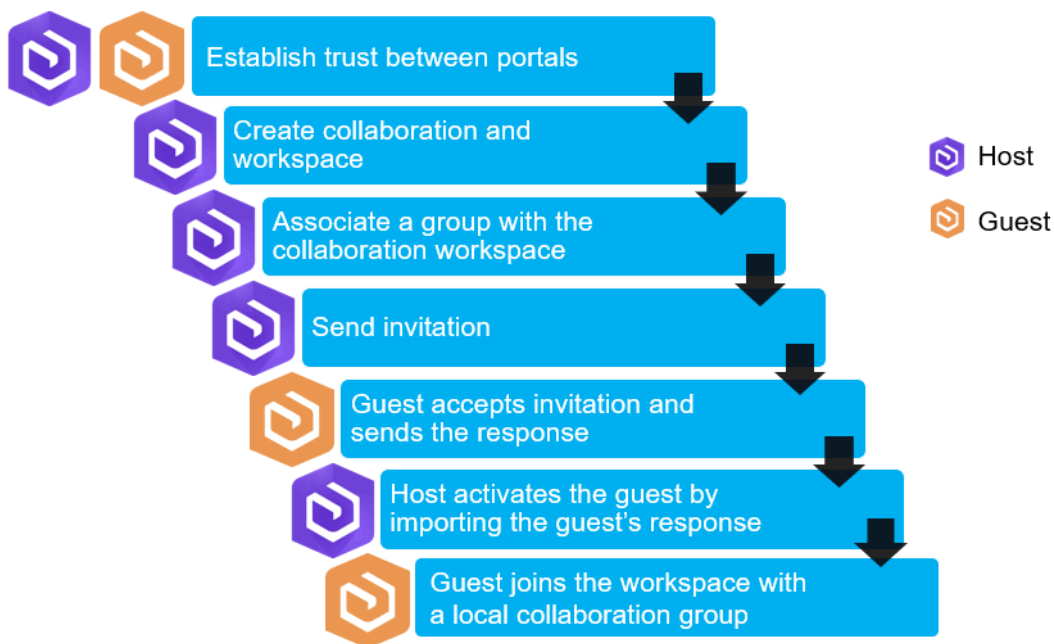


Figure 11.1. Administrators of portals participating in a collaboration must perform these steps to create a new collaboration. The actions that an administrator of the host portal must perform are indicated by the purple ArcGIS Enterprise icon, and the actions that an administrator of guest portals must perform are indicated by the orange ArcGIS Enterprise icon.

1. Administrators will **establish trust between portals**. HTTPS only is a required security setting. Additionally, an ArcGIS Enterprise host must trust the certificates of all guests, and each guest must trust the certificate of the host.
2. The host will **create the collaboration and the workspace** from portal.
3. The host will either **create or use an existing group to associate with the collaboration workspace**.

Steps to create a distributed collaboration (continued)

4. The host will then **send an invitation** to the guest portal.
5. After receiving the invitation, the administrator for guest portal **will accept the invitation in their portal**. Upon completing this action, a response file will be created, and the administrator of the guest portal will **send the response file** to the host.
6. Upon receiving the response file, the administrator for the host portal will **import the guest's response file**.
7. Finally, the guest portal will **join the workspace with either a new or an existing group**.

Explore sharing content in a collaboration

After a distributed collaboration has been set up, the host administrator can create one or more workspaces to facilitate sharing content between organizations. The way in which the workspace is configured will depend on the business needs for your shared project or initiative.

You will use ArcGIS Enterprise Help documentation to explore several key data considerations for sharing data.

Instructions

- a** Open a web browser and go to the [Portal for ArcGIS: Share Content With Collaboration Groups](https://enterprise.arcgis.com/en/portal/11.3/use/share-content-with-collaboration-participants.htm) web page (https://enterprise.arcgis.com/en/portal/11.3/use/share-content-with-collaboration-participants.htm).
- b** Review this page to answer the first question in your workbook.
- c** Scroll down and review the Share Layers And Views section to answer the remaining questions.
- d** After you are finished, close the web browser.

1. Which types of items can be shared with other participants in a collaboration?

2. In which two ways can you share a feature layer using a collaboration?

3. Why might an organization choose to share a feature layer as a copy?

Exercise 11

Use a collaboration to share content

Imagine that you manage multiple ArcGIS Enterprise deployments for a municipality. You are collaborating with the Public Works and Environmental Services Division to help the departments create and maintain datasets. You want to use distributive collaboration to share information with members of the two departments to support functionality like two-way editing for feature layers, showing the results of ongoing tree inspections across the city. You will act as an administrator in both portals to support the initiative.

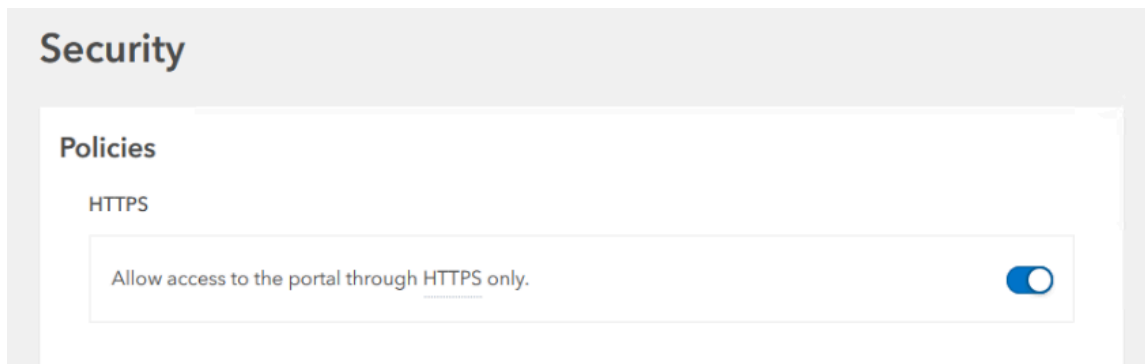
In this exercise, you will perform the following tasks:

- Verify portal security settings.
- Create a collaboration invitation.
- Join a workspace.
- Create a new workspace.

Step 1: Ensure HTTPS only access to each portal

In this step, you will verify both ArcGIS Enterprise portals' security settings to confirm that the portals involved in the collaboration meet the requirements for distributed collaboration.

- a Using Remote Desktop Connection Manager, in the server tree panel, double-click EBASE to connect to the EBASE VM.
- b Open a web browser and click the ArcGIS Enterprise bookmark (<https://ebase.ad.local/portal/home/>) to access the Public Works organization.
- c Sign in to the Enterprise portal as **portaladmin** (password: **Esri.4.GIS**).
- d Click Organization, and then click the Settings tab.
- e On the left, click the Security tab.
- f Notice that the Allow Access To The Portal Through HTTPS Only option is enabled.



For portals to participate in a collaboration, the HTTPS Only option must be enabled.

- g Add a new browser tab, and then click the ArcGIS Enterprise-2 bookmark (<https://ebase2.ad.local/portal/home/>).
- h Sign in to the Enterprise 2 portal as **portaladmin2** (password: **Esri.4.GIS**).
- i Confirm that the Enterprise 2 portal for the Environmental Services Division organization also has the HTTPS only option enabled.

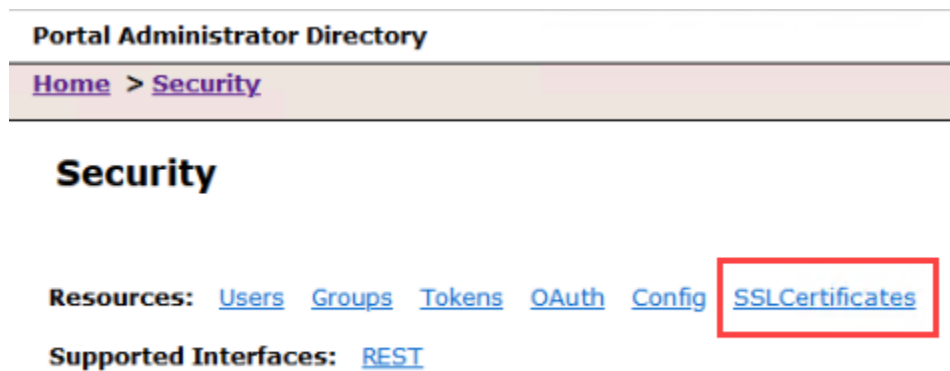
Hint: Organization > Settings > Security

In each organization's portal, you verified that the HTTPS Only option has been enabled.

Step 2: Review the trusted certificates

Next, you will review the trusted certificates in each ArcGIS Enterprise Portal Administrator Site Directory. It is crucial for each ArcGIS Enterprise within the collaboration to trust each other. If the guest organization's portal does not trust the certificate of the host or vice versa, the collaboration will not be successful.

- a Open a new web browser tab and go to the Portal Administrator Site Directory for the Public Works organization (<https://ebase.ad.local/portal/portaladmin/>).
- b Sign in to the Portal Administrator Directory as **portaladmin** (password: **Esri.4.GIS**).
- c Click Security, and then click SSLCertificates, as indicated in the following graphic.



- d Repeat the previous steps for the Environmental Services Division organization.

Hint:

- Open a new web browser tab and go to the Portal Administrator Site Directory for the Environmental Services Division organization (<https://ebase2.ad.local/portal/portaladmin/>).
- Sign in to the Portal Administrator Directory as **portaladmin2** (password: **Esri.4.GIS**).

- e For both portals, click the Adroot certificate to inspect the certificate.

Portal Administrator Directory

[Home](#) > [Security](#) > [SSL Certificates](#) > [adroot](#)

Certificate Information

```

Alias Name: adroot
Entry Type: trustedCertEntry
Issuer: CN=ad-ADSRV2019-CA, DC=ad, DC=local
Subject: CN=ad-ADSRV2019-CA, DC=ad, DC=local
Subject Alternative Names:

Valid From: Wed Mar 09 18:13:33 UTC 2022
Valid Until: Tue Mar 09 18:23:32 UTC 2032
Key Algorithm: RSA
Key Size: 2048
Serial Number: 456be2a96bd9159642f2eb42a0d830ac
Version: 3
Signature Algorithm: SHA256withRSA
Key Usage:
  • Digital_Signature
  • Key_Cert_Sign
  • CRL_Sign

MD5 Fingerprint: 8cc697123be4a616d94278030e3f37de
SHA1 Fingerprint: 81381451032aecb4497fa565d938304f52ce55a0
SHA256 Fingerprint: 015fc21ff1ea7ad5d6361eca4659ef8c63fd289dc97d7d36f314fedc785acfa2

```

Supported Operations: [GenerateCSR](#) [Export](#) [Delete](#) [Import Signed Certificate](#)

Supported Interfaces: [REST](#)

If the host and the guest portals do not have a widely recognized CA certificate, then the root certificate must be imported into the Portal Administrator Site Directory to indicate that the certificate is trusted. In this case, the certificate is the same for both portals, so no additional work will need to be completed.

 Close the Portal Administrator Directory web browser tabs.

In this step, you confirmed that trust has been established between the Public Works and Environmental Services Division portals. Next, you will create the invitation from the Public Works portal, as they will be the host in the collaboration.

Step 3: Create the collaboration invitation

In this step, you will create the invitation from the host, Public Works, to send to the Environmental Services Division guest portal.

- a Restore the Public Works portal (<https://ebase.ad.local/portal/home>) web browser tab that has the portal security settings open.



If you accidentally closed the Public Works portal web browser tab, perform the following steps:

- Click the ArcGIS Enterprise bookmark to access the Public Works organization.
 - Sign in to the Enterprise portal as **portaladmin** (password: **Esri.4.GIS**).
 - Click Organization, and then click the Settings tab.
- b On the left, click the Collaborations tab, and then under Collaborations, click Create A Distributed Collaboration.
 - c In the Create A Distributed Collaboration dialog box, specify the following information:
 - For Collaboration Name, type **Public Works and Environmental Services Division**.
 - For Collaboration Description, type **This collaboration is intended to share information with members of the Public Works Department and the Environmental Services Division to support the results of ongoing projects and initiatives.**
 - d Click Next, and then specify the following information:
 - For Workspace Name, type **Tree Inspections**.
 - For Workspace Description, type **This workspace is to support two-way editing of ongoing tree inspections across the city.**
 - e Click Next, and then specify the following information:
 - Confirm that New Group is selected, and then type **Public Works - Tree Inspections**.
 - For Tags, type **tree inspections, tree survey**.

Because you selected New Group, this action will create a new group to be used for Public Works in the collaboration.

- f** Click Next.
- g** For Feature Layers And Views In My Portal Are Sent As, select Copies.
- h** Check the box for Allow Two-Way Sharing Of Feature Service Edits To Eligible Participants.
- i** Uncheck the box for If Unable To Share As Copies Share As References.

Create a distributed collaboration ✕

Step 4 of 4

Set the workspace sync settings.

Sync Settings Sync immediately

Feature layers and views in my portal are sent as

References

Copies

Collaboration participants receive feature layers and views as hosted feature layers that are periodically synced. Ensure that you check 'Enable Sync' on the source feature layer's or view's item settings to send it as a copy. [Learn more](#)

Allow two-way sharing of feature service edits to eligible participants

If unable to share as copies share as references

Back
Save
Save and Invite
Cancel

This action will allow both departments to edit the feature layer data shared to the Tree Inspections workspace and have those edits synced back to each portal.

- j** Click Save And Invite.

The Invite Guest Organization window opens.

- k** For Guest Organization URL, type **https://ebase2.ad.local/portal/**.

This URL is the portal URL for the guest portal. In this case, the guest portal is the Environmental

Services Division.

- l Under Guest Organization Access To Workspaces, to the right of Tree Inspections, next to Send Content, click the down arrow and choose Send And Receive Content.

To support two-way editing, you will choose Send And Receive Content for the access mode.

- m Under Definitions, notice the access mode definitions, as indicated in the following graphic.

Invite Guest Organization ×

Invite a guest organization to join your collaboration by creating an invitation file.

Guest Organization URL:

Example ArcGIS Enterprise: <https://webadaptorhost.domain.com/webadaptorname/>

Guest Organization Access to Workspaces:

Tree Inspections Send and Receive Content ▼

Definitions:

Send Content: The guest's items will be sent to this workspace.

Receive Content: The guest will receive items from this workspace.

Send and Receive Content: The guest will send items to and receive items from this workspace.

Send Invitation Cancel

- n Click Send Invitation.

An invitation is created and added to the Downloads folder on the host VM.

Typically, the GIS administrator would then send or provide the guest portals access to the invitation. Because you are the GIS administrator overseeing both portals involved in the collaboration, you will copy the invitation from the host machine to the guest machine.

In this step, you began the process of creating a distributed collaboration in the host portal by

creating and copying the collaboration invitation.

Step 4: Accept the collaboration invitation

In this step, you will join the Environmental Services Division portal to the collaboration as the guest by accepting the invitation.

- a Restore the Environmental Services Division's portal (<https://ebase2.ad.local/portal/home>) web browser tab.



If you accidentally closed the Environmental Services Division's portal web browser tab, perform the following steps:

- Click the ArcGIS Enterprise-2 bookmark to access the Environmental Services Division organization.
- Sign in to the Enterprise 2 portal as **portaladmin2** (password: **Esri.4.GIS**).
- Click Organization, and then click the Settings tab.

- b On the left, click the Collaborations tab.
- c Under Collaborations, click Accept Invitation.
- d In the Accept Collaboration Invitation window, click Choose File.
- e Browse to your Downloads folder, select the Collaboration-ebase2.ad.local.invite file, and click Open.
- f In the Accept Collaboration Invitation window, click Accept Invitation, and then click Save Response.

The response is created and saved to your Downloads folder. The Environmental Services Division guest portal created a response that will now need to be received by the Public Works host portal. Similar to the invitation, it is typical for the administrator of the guest portal to send or provide the host portal access to the response file.

- g Restore the Public Works portal (<https://ebase.ad.local/portal/home>) web browser tab.

Collaborations / Public Works and Environmental Services Division: Guests

This collaboration is intended to share information with members of the Public Works Department and the Environmental Services Division to support the results of ongoing projects and initiatives.

Host Organization: ArcGIS Enterprise

Contact Person: portaladmin built-in [EDIT](#)

[Invite Guest](#)

[View Workspaces](#)

| Guest Organization | URL | Email | Status | Action |
|--------------------|----------------------------|-------|--------------------|--------|
| | https://ebase2.ad.local... | | Invitation Pending | |

The status indicates that the invitation from the host portal has been created and is pending a response. You will now receive the response that was created by the guest portal.

- h** Under Action, click the Actions button and choose Accept Guest Organization.
- i** In the Accept Guest Organization window, click Choose File.
- j** Browse to your Downloads folder, select the Collaboration-ebase2.ad.local.response file, and click Open.
- k** In the Accept Guest Organization window, click Accept Guest Organization.

Collaborations / Public Works and Environmental Services Division: Guests

This collaboration is intended to share information with members of the Public Works Department and the Environmental Services Division to support the results of ongoing projects and initiatives.

Host Organization: ArcGIS Enterprise

Contact Person: portaladmin built-in [EDIT](#)

[Invite Guest](#)

[View Workspaces](#)

| Guest Organization | URL | Email | Status | Action |
|---------------------|----------------------------|-----------------------|--------|--------|
| ArcGIS Enterprise 2 | https://ebase2.ad.local... | portaladmin2@ad.local | Active | |

Now the status of the collaboration has changed from Invitation Pending to Active. In this step, you accepted the invitation and generated a response as the guest portal. As the host, you received the response file from the guest; however, there are still steps that need to be completed for the collaboration. The guest portal will need to create or select a group to join

the workspace.

Step 5: Join a workspace from a guest portal

In this step, you will navigate to the Environmental Services Division guest portal to create a new group to join the Tree Inspections workspace.

- a Restore the Environmental Services Division portal (<https://ebase2.ad.local/portal/home>) web browser tab.
- b Under Collaboration Name, click the Public Works And Environmental Services Division collaboration.

Collaborations / Public Works and Environmental Services Division: Workspaces

This collaboration is intended to share information with members of the Public Works Department and the Environmental Services Division to support the results of ongoing projects and initiatives.

Host Organization: ArcGIS Enterprise
 Contact Person: portaladmin built-in
 Requires Web Tier Authentication: No [EDIT](#)

[View Guests](#)

| Workspace | Group | Access | Status | Action |
|------------------|-------|-------------------------|----------------|--------|
| Tree Inspections | | Send and Receive Con... | Not Yet Joined | |

1. What is the status of the Tree Inspections workspace?

Although the guest portal has received the invitation and created the response, it has not joined the workspace. You must join the workspace and associate a group with the workspace to share and receive the tree survey data.

- c For the Tree Inspections workspace, click the Actions button and choose Join Workspace.
- d In the Join Workspace dialog box, specify the following information:
 - For New Group, type **Environmental Services Division - Tree Inspections**.
 - For Tags, type **tree inspections, tree survey**.

- Under Feature Layers And Views In My Portal Are Sent As, select Copies
- Check the box for Allow Two-Way Sharing Of Feature Service Edits To Eligible Participants.
- Uncheck the box for If Unable To Share As Copies Share As References.

Group

New Group:

Tags:
Add tags

Existing Group:

Feature layers and views in my portal are sent as

References

Copies

Collaboration participants receive feature layers and views as hosted feature layers that are periodically synced. Ensure that you check 'Enable Sync' on the source feature layer's or view's item settings to send it as a copy. [Learn more](#)

Allow two-way sharing of feature service edits to eligible participants

If unable to share as copies share as references

- e** Scroll down and locate the Feature Layer Sync Settings section.

As the guest, you are able to configure the sync schedule.

- f** Click Join Workspace, and then click OK to close the message that you successfully joined the workspace.

Acting as an administrator in the guest portal, you joined the new workspace. During this process, you also linked a portal group to the workspaces on both the host and guest portals. At this point, you would typically add portal members to the appropriate collaboration groups.

Step 6: Create a new workspace

Using the same collaboration as before, you will create an additional workspace to support another initiative for the Public Works and Environmental Services Division departments. This

workspace will be used to share information for a different project that the organizations are collaborating on together.

- a Restore the Public Works portal (<https://ebase.ad.local/portal/home>) web browser tab.
- b Click View Workspaces.

The screenshot shows the 'Collaborations' page in ArcGIS. The main heading is 'Collaborations'. Below it, there is a section titled 'Distributed' with explanatory text: 'Distributed collaborations allow ArcGIS organizations to share and sync content using groups. When you create a collaboration, your organization becomes the host. Your organization may also be invited to join a collaboration as a guest. Distributed collaborations are trusted relationships where the terms and conditions set by all organizations apply.'

Below the text, there is a breadcrumb trail: 'Collaborations / Public Works and Environmental Services Division: Workspaces'. A description follows: 'This collaboration is intended to share information with members of the Public Works Department and the Environmental Services Division to support the results of ongoing projects and initiatives.'

Key information includes:

- Host Organization:** ArcGIS Enterprise
- Contact Person:** portaladmin built-in [EDIT](#)

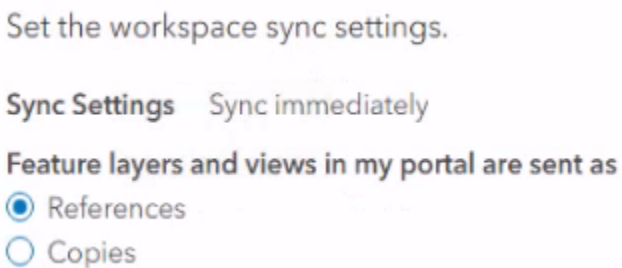
There are two buttons: 'Create Workspace' and 'View Guests'. Below these is a table listing workspaces:

| Workspace | Group | Action |
|------------------|--------------------------------|--------|
| Tree Inspections | Public Works -Tree Inspections | |

- c Click Create Workspace.
- d In the Create Workspace dialog box, specify the following information:
 - For Workspace Name, type **Environmental Assets**.
 - For Workspace Description, type **Public Works department to collaborate with the Environmental Services Division department to update environmental asset data**.
- e Click Next.

In Step 2 of the Create Workspace dialog box, you have the opportunity to specify which group in your organization will be linked to this workspace. For convenience, you can create a group here.

- f In the Create Workspace dialog box (Step 2), specify the following information:
 - For New Group, type **Public Works - Environmental Assets**.
 - For Tags, type **ecosystem services**.
- g Click Next.
- h Under Feature Layers And Views In My Portal Are Sent As, verify that References is selected.



- i Click Next.

You will now define how the Environmental Services department, which is the guest organization, will be able to interact with the workspace. In this case, you want the Environmental Services department to receive content only. The Environmental Services department will not be able to send content to the Public Works department through this workspace.

- j Under Set Each Guest Organization's Workspace Access, to the right of ArcGIS Enterprise 2, click the down arrow, and choose Receive Content.
- k Click Save, and then click OK to close the message that the workspace was created.
- l Under the Workspace heading, click the Environmental Assets link.

To receive content from the Public Works portal, the Environment Service Division would need to join the new workspace using the same steps that you performed earlier. Because new workspace creation only syncs every 15 minutes, it may take that long for this workspace to appear. For the purposes of this exercise, you do not need to complete the steps to join this workspace.

- m Close the browser window.

In this exercise, you acted as the administrator for both the host and the guest portal,

performing all the required steps for both organizations to have a successful collaboration.

Lesson review

1. A city (as the host) wants to create a collaboration workspace to get data updates from a contractor (as the guest), but it does not need to send data to the contractor. Which access mode should the city choose when creating the workspace?
 - a. Disable two-way editing
 - b. Send only
 - c. Receive only
 - d. Send and receive

2. A GIS administrator has an existing distributed collaboration set up between the Parks and Recreation department and the Utilities department for an existing project. The two entities have begun a new project that also requires sharing information between portals. Which action should the GIS administrator perform to support collaboration workflows for both projects?
 - a. Create a new collaboration
 - b. Create a new group
 - c. Create a new workspace
 - d. Adjust the access mode

3. Each participant in a collaboration links a group from their organization to the workspace.
 - a. True
 - b. False

Answers to lesson 11 questions

Explore sharing content in a collaboration (page 11-9)

1. Which types of items can be shared with other participants in a collaboration?

Examples can include web apps, maps, scenes, layers and views, imagery layers, files, and ArcGIS Insights items.

2. In which two ways can you share a feature layer using a collaboration?

Feature layers can be shared by reference or shared as copies.

3. Why might an organization choose to share a feature layer as a copy?

To support two-way editing; both the sending and receiving portal can edit and synchronize updates.

Exercise 11: Use a collaboration to share content (page 11-10)

1. What is the status of the Tree Inspections workspace?

Not Yet Joined

Throughout this course, you have learned about various tasks and workflows that are associated with ArcGIS Enterprise. You may now want to consider automating these workflows for benefits such as mitigating user error, collaboration, and scheduling. In addition to administrative tasks, organizations may also be interested in using automation when conducting data science workflows or management of data in ArcGIS Enterprise. In this lesson, you will review and learn to use automation strategies that can be used to automate administrative tasks that help keep your organization running.

Topics covered

- Automation strategies

- Configuring webhooks

- Using the ArcGIS API for Python to automate workflows

- Automating notebook execution

Automating administrative tasks

As an ArcGIS Enterprise administrator, automating routine tasks can introduce a wide range of benefits:

- **Efficiency:** Automation saves you time.
- **Repeatability:** Tasks can be performed in exactly the same way every time.
- **Ease of use:** Automation methods can serve as documentation for business processes.
- **Scheduling:** Tasks can be completed outside of business hours.
- **Collaboration:** Collaborate on automation scripts with colleagues.

What types of ArcGIS Enterprise administration tasks might benefit from automation?

Automation strategies

There are numerous strategies that organizations can use to automate workflows in ArcGIS Enterprise. Before determining a strategy for a task, it is important to first assess the business need and your familiarity with each strategy to determine which would be most efficient and effective. The following table lists and describes the various automation strategies for ArcGIS Enterprise.

| Automation strategy | Description |
|--------------------------|---|
| Configuration management | Uses infrastructure as code to automate ArcGIS Enterprise deployment and configuration with PowerShell DSC and Chef |
| ArcGIS REST API | Uses any language that can make an HTTP request to interact with ArcGIS Enterprise's REST endpoints |
| Shell scripts | Automates terminal commands and invokes ArcGIS Enterprise command line utilities |
| ArcGIS API for Python | Uses Python modules, classes, and methods to interact with ArcGIS Enterprise |
| ArcGIS Notebooks | Development environment that is integrated with ArcGIS for writing and running automation scripts using Python |
| Webhooks | When an event occurs, information is sent to a receiver to produce a desired action |

ArcGIS Notebooks and ArcGIS Notebook Server

The ArcGIS Notebook Server is not required to perform automation in ArcGIS Enterprise; however, the server licensing role makes it easier to implement many automation strategies across your organization. ArcGIS Notebook Server is a complete data science platform integrated with the ArcGIS Enterprise portal that hosts and runs ArcGIS Notebooks, a web-based scripting interface. When using the notebook editor provided by ArcGIS Notebooks, you can use the ArcGIS API for Python and ArcPy, Esri's Python resources, to perform spatial analysis, craft data science and machine learning workflows, manage GIS data and content, and automate administrative tasks. ArcGIS Enterprise comes with a standard license of ArcGIS Notebook Server at no additional cost.

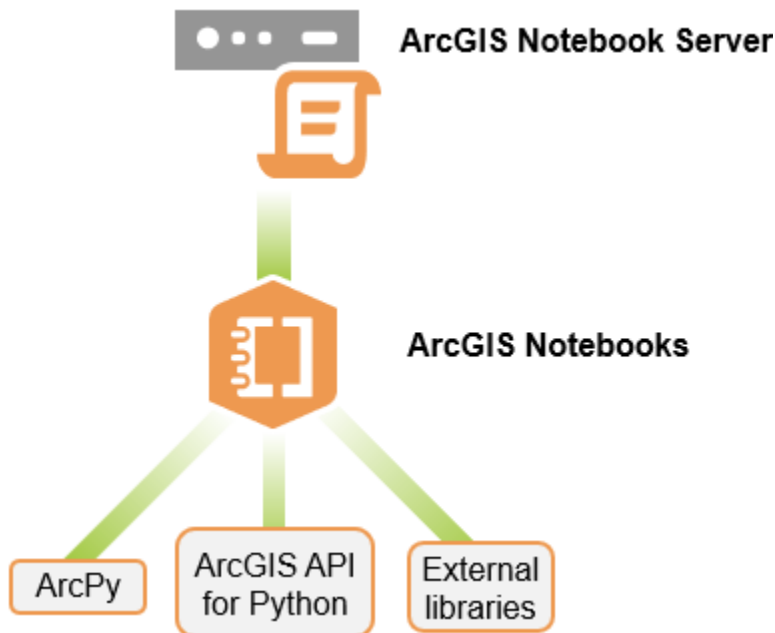


Figure 12.1. ArcGIS Notebook Server allows you to host ArcGIS Notebooks in your Enterprise portal. Using ArcGIS Notebooks, you can use the ArcGIS API for Python, as well as ArcPy and numerous external libraries, to work with your ArcGIS Enterprise portal.

Characteristics of ArcGIS Notebooks in ArcGIS Enterprise

Notebook Server allows you and members of your Enterprise portal to author, document, run, and share Python code using ArcGIS Notebooks. A notebook appears as an item in the portal, just like a layer or web app, and it uses the same identity-based security model to determine

ArcGIS Notebooks and ArcGIS Notebook Server (continued)

access. A notebook allows you to easily work with the content, users, and capabilities in your ArcGIS Enterprise deployment.

| Category | Characteristic |
|-----------------------------|---|
| Working with portal content | Because ArcGIS Notebooks is completely integrated with your portal, searching for content and adding content to a notebook are simple. After content is added to a notebook, you can use the full power of the ArcGIS API for Python to work with it. |
| Analysis and visualization | You can take advantage of your portal's analysis tools, ArcPy, and popular open-source Python modules to analyze and visualize your data. Notebooks allow you to perform analysis and see the results in the same location. |
| Sharing | Notebooks can be easily shared across your ArcGIS Enterprise deployment, allowing authors to share not only their code but also entire workflows and documentation. |

Exercise 12A

Create webhooks as an automation strategy in ArcGIS Enterprise

Imagine that a new analyst has recently joined an organization's GIS team, and the administrator would like to ensure that the team member is added to any new groups that are created in the portal for oversight on various projects. Instead of having to remember to review newly created groups and add the new member manually, the administrator will take advantage of webhooks to automatically add the user to any new group. Because the organization has other automation needs that would benefit from notebooks, the administrator has decided to use ArcGIS Notebook Server to enable the use of notebooks as webhook receivers. They will first test the workflow with an existing member in the organization.

In this exercise, you will perform the following tasks:

- Federate an ArcGIS Notebook Server site.
- Create a custom role.
- Upload a notebook.
- Create and test a webhook.

Step 1: Federate an ArcGIS Notebook Server site

For the purposes of this exercise, you will need to start the machine that is running ArcGIS Notebook Server by running a script.

- a On the host machine, in File Explorer, browse **C:_vm_Scripts** and double-click `Begin_Lesson-12.bat`.
- b Wait for the script to complete.

The `Begin_Lesson-12.bat` script shuts down the EBASE2, GEO, and STBDS VMs and starts the NBSLNX VM. Therefore, the ADSRV2019, EBASE, and NBSLNX VMs should be running for this exercise.

In this step, you will federate an ArcGIS Notebook Server site with your Enterprise portal. The webhook that you will create will use a notebook as a receiver.

- c Using Remote Desktop Connection Manager, in the server tree panel, double-click EBASE to connect to the EBASE VM.
- d In File Explorer, browse to `\\EADM\EsrITraining\EADM\AutomationMgmt` and open the `FederateNotebookServer.txt` file.

This text file contains the URLs that you will copy to ensure accuracy as you federate the server.

- e Open a web browser and click the ArcGIS Enterprise bookmark (<https://ebase.ad.local/portal/home/>).
- f Sign in to the Enterprise portal as **portaladmin** (password: **Esri.4.GIS**).
- g Click Organization, if necessary, and then click the Settings tab to access the ArcGIS Enterprise settings page.
- h On the left side, click the Servers tab.
- i Click Add Server Site, and then specify the following information:
 - Services URL: **https://ebase.ad.local/notebk_server** (copy and paste from the text file)
 - Administration URL: **<https://nbslnx.ad.local:11443/arcgis>** (copy and paste from the text file)

file)

- Username: **siteadmin**
- Password: **Esri.4.GIS**

1. Why do the Services URL and the Administration URL reference different machine names?

In this case, the ArcGIS Web Adaptor instance that forwards requests to the Notebook Server site is installed on the EBASE machine. This setup mimics a common scenario in which web adaptors are installed on web servers rather than the server host.

- j** Click Next.
- k** Enable the Notebook Server role.
- l** Click Save Server Role.

The list of federated servers will update to include your Notebook Server site. Because you are signed in as an administrator, you have immediate access to the Notebook tab.

- m** Refresh the page, if necessary, and then verify that the Notebook tab is visible near the top of the page.
- n** Click Notebook.

From the Notebook tab, you can access all your notebooks, notebooks that have been shared with you, and Esri sample notebooks. Additionally, you can manage notebook environments.

In this step, you confirmed that your Notebook Server site is federated, and you are ready to allow others within your organization to access its capabilities.

Step 2: Create a custom role for notebook authors

The next step is to create a role within your Enterprise portal for members who must author notebooks.

- a Click Organization, and then click the Settings tab.
- b On the left side, click the Member Roles tab.
- c Click Create Role.
- d In the Create Role dialog box, for Role Name, type **Notebook Author**.
- e For Description, type **A role that includes all notebook-related privileges**.

You will now specify the privileges associated with this role, keeping in mind that you want this role to include all the privileges of the default Publisher role.

- f In the Role Privileges section, click Set From Existing Role.
- g In the Import Settings From An Existing Role Or Template window, for Select Role Or Template To Import Settings From, choose Publisher.
- h Click Import Settings.
- i Under General Privileges, expand the Content category.
- j In the Content category, locate the following privileges:
 - Create And Edit Notebooks
 - Schedule Notebooks

2. Does the default Publisher role have these privileges enabled?

- k Enable the Create And Edit Notebooks and Schedule Notebooks privileges.

Create and edit notebooks

Allow member to create and edit interactive notebooks.



Schedule notebooks

Allow member to schedule future automated runs of a notebook.



The Create And Edit Notebooks privilege allows members to create notebooks and edit their existing notebooks in the portal. The Schedule Notebooks privilege allows members to schedule when a notebook is run.

- l Expand the Content And Analysis category, and then enable the Advanced Notebooks privilege.



To use ArcPy in ArcGIS Notebooks, the Advanced Notebooks privilege must be assigned, along with the Create And Edit Notebooks privilege.

- m Click Save to create the new role.
- n Near the top of the page, click the Members tab.
- o For the Publisher Built-in member, in the Role column, click the down arrow and choose Notebook Author.

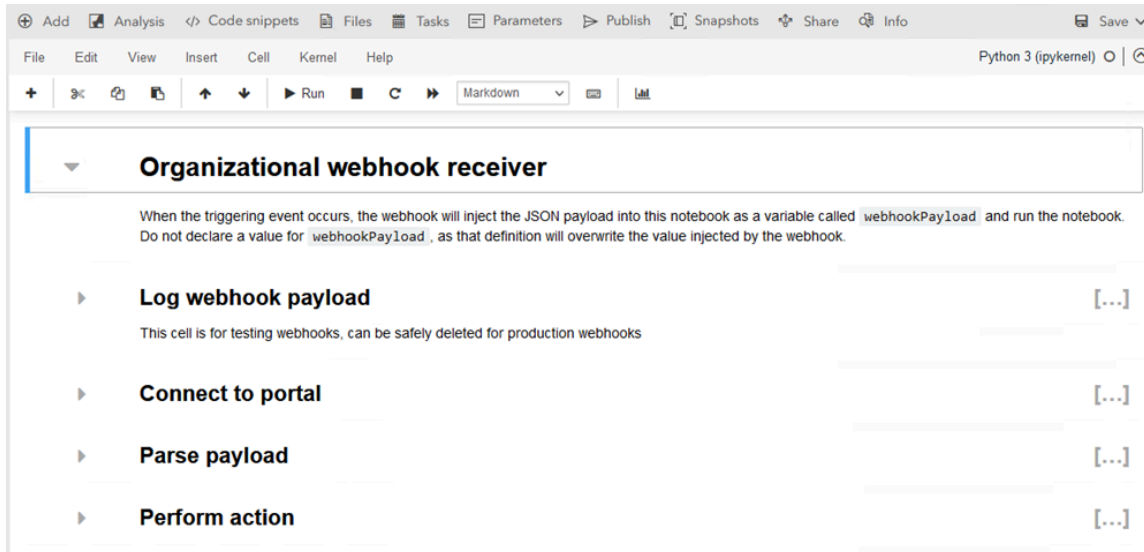


Step 3: Upload a notebook

Next, you will upload an existing notebook file, which will create a portal item in your Enterprise portal.

- a Click Content.
- b Click New Item, and then in the New Item window, click Your Device.
- c Browse to `\\EADM\EsriTraining\EADM\AutomationMgmt.`
- d Click the `AddMemberToGroupsWebhook.ipynb` file to select it, and then click Open.
- e For Tags, type **webhook, group members**.
- f For Select A Runtime To Be Used With Your Notebook, choose ArcGIS Notebook Python 3 Advanced - 10.0.

- g** Click Save to upload the notebook to the portal.
- h** On the item page, click Open Notebook to review the notebook.
- i** Review the contents of the notebook and notice the four sections that are included in the notebook.



The sections in the notebook were collapsed in this screenshot.

This notebook will be used as the receiver for the webhook that you will create. In this scenario, when a new group is created in ArcGIS Enterprise, it will trigger the webhook. The webhook will then create and capture information in what is referred to as a payload. The payload is injected into the notebook. Using the injected information, the notebook executes, adding the publisher member to the group.

- j** After you have finished reviewing the notebook, close the web browser tab and, if necessary, choose to leave the page.

Step 4: Create a webhook

Now that the script is uploaded as an item in the portal, you will create a webhook that will trigger when a new group is created.

- a** Click Organization.
- b** Click the Settings tab, and then on the left side, click the Webhooks tab.

- c Click Create Webhook.
- d In the Create Webhook window, specify the following information:
 - For Name, type **Add user to new group**.
 - For Type, choose Groups.
 - For Groups, verify that All Groups is selected.
 - For Events, choose /groups/add.
 - Select Notebook, and then from the drop-down list, choose AddMemberToGroupsWebhook

The screenshot shows the 'Create webhook' dialog box. The fields are filled as follows:

- Name:** Add user to new group
- Type:** Groups
- Groups:** All groups
- Events:** /groups/add (Add trigger events)
- Webhook URL:** (unselected)
- Notebook:** (selected)
- Notebook Name:** AddMemberToGroupsWebhook
- Secret (optional):** Add secret key

Buttons at the bottom: Cancel, Create webhook

You will not choose Secret because you do not need to authenticate payload information with the receiver.

Because the ArcGIS Notebook Server licensing role is configured with ArcGIS Enterprise, you have the option to select a notebook to be the receiver.

- e Click Create Webhook.

The webhook is now configured to trigger when any new group is added in the portal. When it is triggered, the webhook will send the payload to the selected notebook, AddMemberToGroupsWebhook.

Step 5: Test the webhook

You will now test the webhook by creating a new group.

- a At the top of the page, click Groups.
- b Click Create Group.
- c In the Create A Group window, specify the following information:
 - Name: **TestWebhook**
 - Summary: **Testing webhooks**
 - Tags: **Test**
- d Leave the remaining defaults and click Save.
- e On the group page, click the Members tab.



If you do not see publisher listed as a member, refresh the browser.

Now both the portaladmin and the publisher are members of the group.

- f Close the FederateNotebookServer.txt file and close File Explorer.
- g Leave the web browser open for the next exercise.

In this exercise, the ArcGIS Notebook Server was federated as a solution to easily integrate automation workflows for the organization. You demonstrated this solution by creating a webhook to use ArcGIS Notebooks so that when any new group is added, the publisher member is added to the group. Although webhooks can be configured without ArcGIS Notebook Server, the additional licensing role provided the ability to configure a webhook to use an ArcGIS Notebooks item as a receiver.

Working with the ArcGIS API for Python

Although notebooks are not required to use the ArcGIS API for Python, it is recommended. Because ArcGIS Notebook Server is federated, and is therefore closely integrated with the ArcGIS Enterprise organization, the authoring experience in a notebook has several unique characteristics that can enhance your ability to efficiently automate workflows for your ArcGIS Enterprise deployment.

The notebook provides the ability to easily search content and analysis tools that are accessible from the portal. ArcGIS API for Python syntax can be directly added to the notebook when accessing items and analysis from ArcGIS Enterprise, making scripting easier. The development environment uses the ArcGIS API for Python to automate workflows such as user and item management, working with web services, and conducting analysis.

Working with the ArcGIS API for Python (continued)

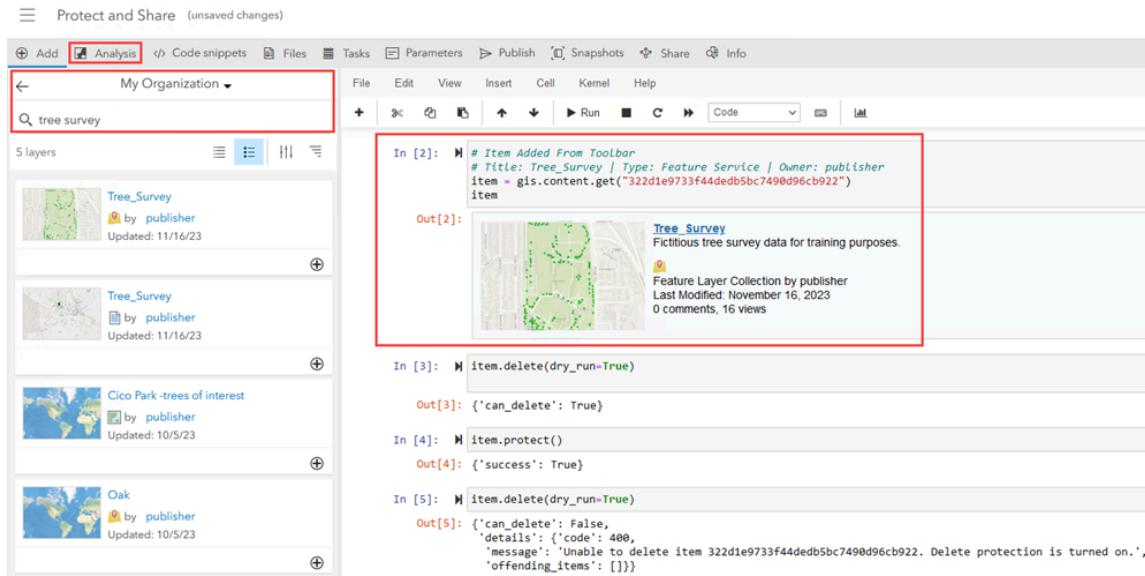


Figure 12.2. In ArcGIS Notebooks, the notebook editor interface allows you to search for and add content, as highlighted in this graphic.

Architecture considerations for ArcGIS Notebook Server

ArcGIS Notebook Server is an ArcGIS Server role, so it has some architectural similarities with other server licensing roles. However, Notebook Server has some special considerations that should be accounted for when deploying a new ArcGIS Notebook Server site.

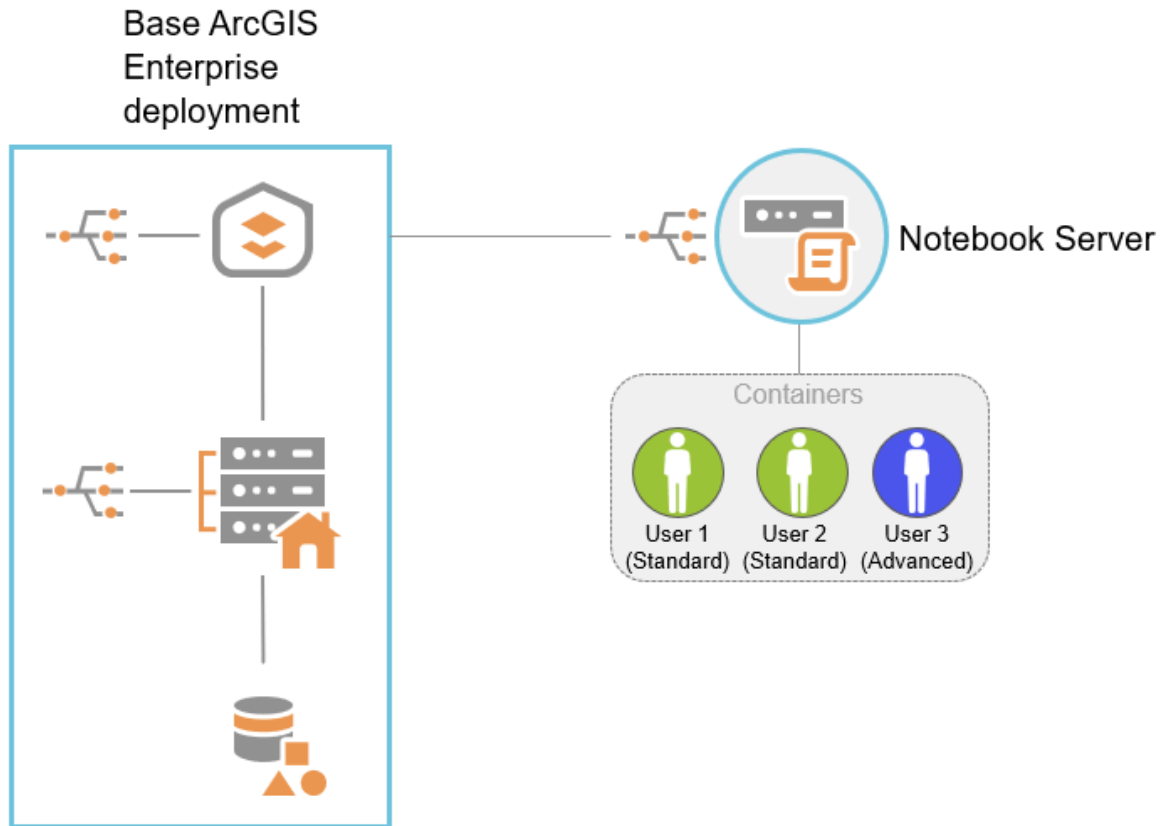


Figure 12.3. In this example, an ArcGIS Notebook Server site is federated with a base ArcGIS Enterprise deployment. The site manages containers that allow notebook authors to work in isolated environments with dedicated resources.

If you are planning to deploy an ArcGIS Notebook Server site, it is important to understand the following architectural considerations:

Architecture considerations for ArcGIS Notebook Server (continued)

| Consideration | Description |
|-------------------|--|
| Federation | <ul style="list-style-type: none"> • Notebook Server is installed on a dedicated server operating system. • Notebook Server must be federated with an Enterprise portal. |
| Containers | <ul style="list-style-type: none"> • Containers are a way to isolate each notebook author's environment. • Containers allow each author to use a limited subset of your server resources. • Your system resources should correspond to the number of notebook authors that you plan to enable. |
| Notebook runtimes | <ul style="list-style-type: none"> • Notebook runtimes are applied to containers through Esri-provided container images. • The Standard runtime includes the ArcGIS API for Python and many other Python modules. • The Advanced runtime includes everything in the Standard runtime, plus ArcPy. |



ArcGIS Enterprise administrators can also limit the number of notebooks that are actively scheduled to ensure that automated tasks do not overwhelm system resources.

Exercise 12B

Use ArcGIS Notebooks to schedule administrative tasks in ArcGIS

You have successfully granted members of your organization the privileges needed to use ArcGIS Notebooks. In this exercise, you will continue exploring the functionality of notebooks by playing the role of a portal member who must schedule a notebook to run once a month. You will also explore how to interact with the user workspace available to each notebook author, as well as the capabilities for managing containers.

In this exercise, you will perform the following tasks:

- Schedule a notebook task.
- Manage notebook containers.



ArcGIS Enterprise Help: *Manage ArcGIS Notebook Server resources*

Step 1: Import a notebook

You will upload an existing notebook file, which will create a portal item in your Enterprise portal.

- a Using Remote Desktop Connection Manager, in the server tree panel, double-click EBASE to connect to the EBASE VM.
- b Restore the web browser.
- c Sign out of ArcGIS Enterprise, and then sign in as **publisher** (password: **Esri.4.GIS**).
- d Click Content.
- e Click New Item, and then in the New Item window, click Your Device.
- f Browse to `\\EADM\EsrITraining\EADM\AutomationMgmt`, select the `ListNewContent.ipynb` file, and click Open.
- g For Tags, type **Admin task**.
- h For Select A Runtime To Be Used With Your Notebook, choose ArcGIS Notebook Python 3 Advanced - 10.0.



New item

File
ListNewContent.ipynb

Title
ListNewContent

Folder
publisher

Tags
Admin task x Add tags x

Summary
Add a summary
Characters left: 2048

Select a runtime to be used with your notebook:
ArcGIS Notebook Python 3 Advanced - 10.0

- i Leave all remaining defaults and click Save.

ArcGIS Enterprise recognizes the item as a notebook, and a new portal item is created. This item has much of the same core functionality as other portal items, including the ability to share the notebook.

- j On the item page, click Open Notebook.

It will take several seconds to create a new notebook container and load the content that will be displayed in the notebook editor.

Step 2: Run a notebook

You will now run the code in the notebook one cell at a time. One of the features of notebooks is the ability to include explanatory text (markdown) with the code. You are encouraged to read the markdown content as you execute the code.

The first cell contains a brief explanation of the purpose of the notebook.

- a Under the List New Content heading, read the explanation.
- b Scroll to the first cell that contains Python code.
- c Click the first cell to highlight it.

```
In [ ]: # First, import some libraries we'll need, and connect to our GIS
import datetime as dt
import csv
from IPython.core.display import display, HTML
from pathlib import Path

from arcgis.gis import GIS
gis = GIS("https://ebase.ad.local/portal", username="publisher", password='Esri.4.GIS')
gis
```



A green outline indicates that the cell is selected and that you can edit the code.

Using the current user's "home" credentials can be a convenient way to create a connection to a GIS and allows you to avoid including account credentials in a notebook.

- d In the first cell, to the right of the equal sign, replace all the text with GIS ("home").

The last three lines of code in the block should look like the following graphic.

```
from arcgis.gis import GIS
gis = GIS("home")
gis
```

- e To execute the cell, near the top of the notebook editor, click Run.



You can also press Shift+Enter to run the cell.

The first execution will take a few minutes to connect to the portal and populate the large GIS object. While a cell is executing or waiting to execute, it will display an asterisk (*) in the label to the left of the cell. When cell execution completes, a number will replace the asterisk, as shown in the following graphic.

```
In [1]: # First, import some libraries
import datetime as dt
import csv
```

- f Below the code, verify that an output cell that displays the root of the GIS object appears.

```
Out[1]: GIS @ https://ebase.ad.local/portal
```

- g Click the second block of code to select it.

```
In [ ]: now = dt.datetime.now()
daysToQuery = 30
start_time = now - dt.timedelta(days=daysToQuery)

# We need to convert seconds to milliseconds to search portal content
now_milliseconds = now.timestamp() * 1000
start_time_milliseconds = start_time.timestamp() * 1000

# This function formats a date & time from the timestamps that we'll receive from the
def format_datetime(age_timestamp):
    datetime_object = dt.datetime.fromtimestamp(age_timestamp / 1000)
    return datetime_object.strftime("%B %d, %Y %I:%M %p")

# The search uses timestamps (in milliseconds)
search_query = "created:{{}} TO {{}}".format(start_time_milliseconds, now_milliseconds)
print('This is the query we will pass into the search: {}'.format(search_query),
```

The code in this cell will construct a search query that you will use to search for recently created content in your Enterprise portal. The `daysToQuery` variable specifies how many days to look back to find results. For testing purposes, you will set this variable to a very large number.

- h Change the `daysToQuery` variable from 30 to 999.

Hint: The second line in the cell should be `daysToQuery = 999`.

- i Run the cell.
- j Read the information generated below the cell.
- k Under the Find And Display The New Content heading, click the third block of code to select it.

```
In [ ]: new_content = gis.content.search(query=search_query, max_items=1000)

# Sort the list to show the most recent items on top
new_content.sort(key= lambda item: item['created'], reverse=True)

for item in new_content:
    date_formatted = format_datetime(item.created)
    print("{} {} {}".format(date_formatted, item.title, item.type, item.owner))
```

- l Run the cell.

Content is listed below the cell, with the most recently created content at the top of the list.

- m Run the fourth block of code.

The result visualizes the same data, but it is formatted as an HTML table. One of the benefits of notebooks is the ability to write code and visualize content in the same environment.

- n Run the fifth block of code, which saves your results to a new CSV file.

```
# Save a csv file to your user workspace

data_path = Path('/arcgis/home/csv_output')
# data_path = Path('C:/EsriTraining/EADM/AutomationMgmt/csv_output')

if not data_path.exists():
    data_path.mkdir()

curTime = dt.datetime.now().strftime('%Y-%m-%d_%H.%M.%S')
file_path = 'new_content.{}.csv'.format(curTime)
save_path = data_path.joinpath(file_path)

with open(save_path, 'w', newline='') as f:
    writer = csv.writer(f)
    writer.writerow(["Publish Date", "Item", "Item Type", "Owner"])
    for item in new_content:
        date_formatted = format_datetime(item.created)
        writer.writerow([date_formatted, item.title, item.type, item.owner])
data_path
```

- o Run the final block of code, which creates a portal item (with current date appended) from the CSV file that you just created.

```
# Create a portal item from the csv
csv_item = gis.content.add({'access':'private'}, str(save_path))
csv_item
```



The results are now accessible as a portal item.

- p** In the top-right corner of the notebook editor, click Save and choose Save to save the edits to the notebook.

You manually ran each cell in the notebook to explore the notebook environment and review the script; however, you can configure notebooks to be executed automatically at a specific time.

Step 3: Schedule a notebook task

Now that you know that the notebook works as expected, you are ready to automate the usage by scheduling the task to run on a monthly basis.

- a** On the top ribbon of the notebook editor, click Tasks.

The Tasks pane opens, where you can create and schedule tasks.

- b** In the Tasks pane, click Create Task.
- c** In the Create Task dialog box, for Title, type **Monthly Content Report**.

Because you do not have any parameters that you want to add into the notebook when it is executed by the task, you can proceed to set the schedule details.

- d** Click Next.
- e** Under Set Task Schedule Details, specify the following parameters:

- Beginning On: <today's date>

- Repeat Type: Month
- Day Of Month: 1
- Repeat Interval: 1 - Every Month
- Time: 01:00 AM
- Ending On: Never

Create task ×

1. Set basic information and parameters

Beginning on

Day of month

Time


Ending on

2. Set task schedule details

Repeat type

Repeat interval

f Leave the remaining parameters as the default values and click Create.

Your task appears in the Tasks pane. You can manage your notebook tasks in the Tasks pane of the notebook editor or through the Manage Tasks page (Menu button  > Notebook > Manage > Manage Tasks).

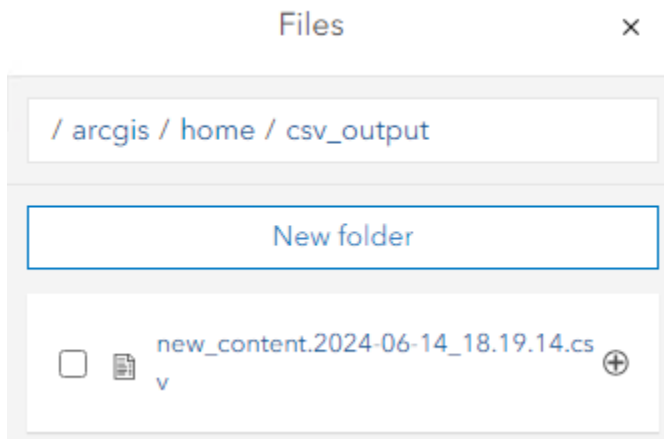
You have now automated the running of an administrative task. You can use the same approach to perform any manner of functions that a member or an administrator may require.

Step 4: Investigate a user workspace

In addition to publishing a CSV as a portal item, this notebook task also saves the CSV file to a private ArcGIS Notebook Server workspace. Each portal member who can create notebooks is provided a private workspace to upload and download files. The user workspace can also be used to set up a scratch workspace to use with ArcPy functions. You will explore your user

workspace next.

- a On the top ribbon of the notebook editor, click Files to open the Files pane.
- b In the Files pane, click the Home link to open the folder.
- c In the Home folder, click the Csv_output link to view the contents of the folder.



The CSV file is saved in your workspace. From the Files pane, you can use the Insert Into Notebook button ⊕ to quickly insert a reference to this file into a notebook.

- d Close the web browser tab showing the notebook editor without saving any changes.

Step 5: Inspect notebook containers

Now you will view active containers as an administrative user using ArcGIS Notebook Server's management tools.

- a Sign out of the Enterprise portal as publisher, and then sign in as **portaladmin** (password: **Esri.4.GIS**).
- b Click Notebook, and then near the upper-left corner, click Manage and choose Manage Notebooks.


| Notebook environment | Version | Type | Notebooks | Member |
|-------------------------------------|---------|-------------|-----------|-------------|
| > ArcGIS Notebook Python 3 Advanced | 10.0 | Interactive | 2 | publisher |
| > ArcGIS Notebook Python 3 Advanced | 10.0 | Interactive | 1 | portaladmin |

There are two active containers. As an administrator, you can see the container associated with yourself as well as the publisher user's container. The publisher user can see their container but not those of other users.

1. What is the runtime used by the publisher's container?

When a container is launched, one of two container images is applied to the container. Because the publisher member has been assigned a role with the Advanced Notebooks privilege, the publisher's container is able to use the advanced container image.

Because containers use computing resources, you may decide to terminate containers that are not in use or not needed. This action will not delete notebook portal items.

- c For the publisher container, click the Delete button .
- d In the Terminate Container window, click Terminate.

When the publisher member opens a notebook again, a new container will be generated.

- e Close the web browser.

In this exercise, you explored how to schedule a task to automate a routine process. You also saw how to access files in the private user workspace and manage notebook containers.

Lesson review

1. A GIS administrator is working with their IT department to efficiently and effectively deploy ArcGIS Enterprise multiple times across the organization to allow different departments to each have their own ArcGIS Enterprise organization. Which automation strategy should the GIS administrator consider to install repeatable ArcGIS Enterprise deployments?
 - a. REST API
 - b. Configuration management
 - c. Webhooks
 - d. ArcGIS API for Python
2. Webhooks require ArcGIS Notebooks for automation workflows in ArcGIS Enterprise.
 - a. True
 - b. False
3. The default Publisher role includes the privileges required for portal members to author and run notebooks.
 - a. True
 - b. False
4. Which ArcGIS Notebook Server component allows each author to use a subset of your server resources without the author's work or resource use interfering with others?
 - a. Workspaces
 - b. Containers
 - c. Runtimes
 - d. Notebooks

Answers to lesson 12 questions

Automating administrative tasks (page 12-2)

What types of ArcGIS Enterprise administration tasks might benefit from automation?

Possible responses include the following:

- **Installation and configuration (for example, mirroring an ArcGIS Enterprise deployment)**
- **User and content management (for example, cloning content in bulk)**
- **Quality control (for example, validating data collected in the field)**
- **Analysis (for example, extracting features from new aerial imagery)**
- **System maintenance (for example, validating your federated servers and registered data stores)**

Exercise 12A: Create webhooks as an automation strategy in ArcGIS Enterprise (page 12-6)

1. Why do the Services URL and the Administration URL reference different machine names?

The web adaptor (notebk_server) is installed on a different machine than Notebook Server.

2. Does the default Publisher role have these privileges enabled?

No; the default Publisher role does not include the ability to create and edit notebooks.

Exercise 12B: Use ArcGIS Notebooks to schedule administrative tasks in ArcGIS Enterprise (page 12-18)

1. What is the runtime used by the publisher's container?

ArcGIS Notebook Python 3 Advanced

Appendix A

Data license agreement

ENVIRONMENTAL SYSTEMS RESEARCH INSTITUTE, INC. ("ESRI"), IS WILLING TO LICENSE THE ENCLOSED ELECTRONIC VERSION OF THE TRAINING MATERIALS TO THE STUDENT ("YOU") ONLY UPON THE CONDITION THAT YOU ACCEPT ALL TERMS AND CONDITIONS CONTAINED IN THIS ESRI DATA LICENSE AGREEMENT ("AGREEMENT"). PLEASE READ THE TERMS AND CONDITIONS CAREFULLY. BY CLICKING, "I ACCEPT", YOU ARE INDICATING YOUR ACCEPTANCE OF THE ESRI DATA LICENSE AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS AND CONDITIONS AS STATED, ESRI IS UNWILLING TO LICENSE THE TRAINING MATERIALS TO YOU.

Training Materials Reservation of Ownership. This Agreement gives You certain limited rights to use electronic and tangible versions of the digital or printed content required to complete a course, which may include, but are not limited to, workbooks, data, concepts, exercises, and exams ("Training Materials"). Esri and its licensor(s) retain exclusive rights, title, and ownership to the copy of Training Materials, software, data, and documentation licensed under this Agreement. Training Materials are protected by United States copyright laws and applicable international copyright treaties and/or conventions. All rights not specifically granted in this Agreement are reserved to Esri and its licensor(s).

Grant of License. Esri grants to You a personal, nonexclusive, nontransferable license to use Training Materials for Your own training purposes. You may run and install one (1) copy of Training Materials and reproduce one (1) copy of Training Materials. You may make one (1) additional copy of the original Training Materials for archive purposes only, unless Esri grants in writing the right to make additional copies.

Training Materials are intended solely for the use of the training of the individual who registered and attended a specific training course. You may not (i) separate the component parts of the Training Materials for use on multiple systems or in the cloud, use in conjunction with any other software package, and/or merge and compile into a separate database(s) or documents for other analytical uses; (ii) make any attempt to circumvent the technological measure(s) (e.g., software or hardware key) that effectively controls access to Training Materials; (iii) remove or obscure any copyright, trademark, and/or proprietary rights notices of Esri or its licensor(s); or (iv) use audio and/or video recording equipment during a training course.

Appendix A

Data license agreement (continued)

Term. The license granted by this Agreement will commence upon Your receipt of the Training Materials and continue until such time that (1) You elect to discontinue use of the Training Materials or (2) Esri terminates this Agreement for Your material breach of this Agreement. This Agreement will be terminated automatically without notice if You fail to comply with any provision of this Agreement. Upon termination of this Agreement in either instance, You will return to Esri or destroy all copies of the Training Materials, including any whole or partial copies in any form, and deliver evidence of such destruction to Esri, and which evidence will be in a form acceptable to Esri in its sole discretion. The parties hereby agree that all provisions that operate to protect the rights of Esri and its licensor(s) will remain in force should breach occur.

Limited Warranty. Esri warrants that the media on which Training Materials is provided will be free from defects in materials and workmanship under normal use and service for a period of ninety (90) days from the date of receipt.

Disclaimer of Warranties. EXCEPT FOR THE LIMITED WARRANTY SET FORTH ABOVE, THE TRAINING AND TRAINING MATERIALS CONTAINED THEREIN ARE PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NONINFRINGEMENT. ESRI DOES NOT WARRANT THAT THE TRAINING OR TRAINING MATERIALS WILL MEET YOUR NEEDS OR EXPECTATIONS; THAT THE USE OF TRAINING MATERIALS WILL BE UNINTERRUPTED; OR THAT ALL NONCONFORMITIES, DEFECTS, OR ERRORS CAN OR WILL BE CORRECTED. THE TRAINING DATABASE HAS BEEN OBTAINED FROM SOURCES BELIEVED TO BE RELIABLE, BUT ITS ACCURACY AND COMPLETENESS, AND THE OPINIONS BASED THEREON, ARE NOT GUARANTEED. THE TRAINING DATABASE MAY CONTAIN SOME NONCONFORMITIES, DEFECTS, ERRORS, AND/OR OMISSIONS. ESRI AND ITS LICENSOR(S) DO NOT WARRANT THAT THE TRAINING DATABASE WILL MEET YOUR NEEDS OR EXPECTATIONS, THAT THE USE OF THE TRAINING DATABASE WILL BE UNINTERRUPTED, OR THAT ALL NONCONFORMITIES CAN OR WILL BE CORRECTED. ESRI AND ITS LICENSOR(S) ARE NOT INVITING RELIANCE ON THIS TRAINING DATABASE, AND YOU SHOULD ALWAYS VERIFY ACTUAL DATA, SUCH AS MAP, SPATIAL, RASTER, OR TABULAR INFORMATION. THE DATA CONTAINED IN THIS PACKAGE IS SUBJECT TO CHANGE WITHOUT NOTICE. IN ADDITION TO AND WITHOUT LIMITING THE PRECEDING PARAGRAPH, ESRI DOES NOT WARRANT IN ANY WAY TRAINING DATA. TRAINING DATA MAY

Appendix A

Data license agreement (continued)

NOT BE FREE OF NONCONFORMITIES, DEFECTS, ERRORS, OR OMISSIONS; BE AVAILABLE WITHOUT INTERRUPTION; BE CORRECTED IF ERRORS ARE DISCOVERED; OR MEET YOUR NEEDS OR EXPECTATIONS. YOU SHOULD NOT RELY ON ANY TRAINING DATA UNLESS YOU HAVE VERIFIED TRAINING DATA AGAINST ACTUAL DATA FROM DOCUMENTS OF RECORD, FIELD MEASUREMENT, OR OBSERVATION.

Exclusive Remedy. Your exclusive remedy and Esri's entire liability for breach of the limited warranties set forth above will be limited, at Esri's sole discretion, to (i) replacement of any defective Training Materials; (ii) repair, correction, or a workaround for Training Materials; or (iii) return of the fees paid by You for Training Material that do not meet Esri's limited warranty, provided that You uninstall, remove, and destroy all copies of the Training Materials and execute and deliver evidence of such actions to Esri.

IN NO EVENT WILL ESRI BE LIABLE TO YOU FOR COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR TRAINING; LOST PROFITS; LOST SALES; BUSINESS EXPENDITURES; INVESTMENTS; BUSINESS COMMITMENTS; LOSS OF ANY GOODWILL; OR ANY INDIRECT, SPECIAL, EXEMPLARY, CONSEQUENTIAL, OR INCIDENTAL DAMAGES ARISING OUT OF OR RELATED TO THIS AGREEMENT, HOWEVER CAUSED OR UNDER ANY THEORY OF LIABILITY, EVEN IF ESRI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. ESRI'S TOTAL CUMULATIVE LIABILITY HEREUNDER, FROM ALL CAUSES OF ACTION OF ANY KIND, WILL IN NO EVENT EXCEED THE AMOUNT ACTUALLY PAID BY YOU FOR THE PORTION OF THE TRAINING UNDER THIS AGREEMENT. THESE LIMITATIONS WILL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY.

Export Regulation. You must comply with all applicable laws and regulations of the United States including, without limitation, its export control laws. You expressly acknowledge and agree not to export, reexport, transfer, or release Esri-provided Training Materials, in whole or in part, to (i) any US embargoed country (including to a resident of any US embargoed country); (ii) any person or entity on the US Treasury Department Specially Designated Nationals List; (iii) any person or entity on the US Commerce Department Lists of Parties of Concern; or (iv) any person or entity where such export, reexport, or provision violates any US export control laws or regulations including, but not limited to, the terms of any export license or licensing provision and any amendments and supplemental additions to US export laws.

Appendix A

Data license agreement (continued)

Governing Law. This Agreement is governed by and construed in accordance with the laws of the state in which training is being held or, in the case of training provided over the Internet, the laws of the State of California, without reference to its conflict of laws principles.

Appendix B

Answers to lesson review questions

Answers to lesson 1 review questions

1. Before implementing your ArcGIS Enterprise deployment, it is crucial to thoroughly assess your business needs.
 - a. **True**
2. What three tiers make up the architecture of an ArcGIS system?
Applications (apps), services, and data
3. Which server licensing role provides capabilities that support the location services ArcGIS system pattern?
ArcGIS GIS Server

Answers to lesson 2 review questions

1. A raster analysis is run using the raster analytics server. Which component is needed to store the output?
 - d. **A raster data store**
2. What capabilities are gained by federating an Image Server site with the ArcGIS Enterprise portal?
The site is fully integrated with your Enterprise portal's security and sharing model and can be used for image hosting and raster analytics.

Appendix B

Answers to lesson review questions (continued)

Answers to lesson 3 review questions

1. High availability is a design approach that allocates functions of your system to specific servers to prevent resource contention.
 - b. **False**
(This statement describes workload separation. High availability is a design approach that minimizes downtime through duplication and load balancing to reduce single points of failure.)
2. Which example is the most representative of workload separation?
 - d. **Dedicating a faster CPU machine for analysis and a weaker machine for visualization**
3. Which best practice is represented by an ArcGIS Enterprise deployment with one primary ArcGIS Data Store and one standby ArcGIS Data Store?
 - c. **High availability**

Answers to lesson 4 review questions

1. True or false: When you share a new web layer from ArcGIS Pro, a new web service is published and exposed as a web layer item in the Enterprise portal.
 - a. **True**
2. Which ArcGIS software component exposes data as services?
 - c. **ArcGIS Server**
3. Which layer utilizes user-managed data?
 - b. **A map image layer referencing an enterprise geodatabase**

Appendix B

Answers to lesson review questions (continued)

Answers to lesson 5 review questions

1. An analyst wants to publish a branch-versioned feature layer where the data is not copied and is maintained in a user-managed data store. Which type of user-managed data store does the analyst need to register?
d. Enterprise geodatabase
2. A publisher has a file geodatabase that contains several rasters and an associated mosaic dataset. The publisher wants to share an imagery layer that references the mosaic dataset. What process should the publisher perform to accomplish this task?
d. Register the folder containing the file geodatabase with an ArcGIS Image Server Site, then publish the imagery layer from ArcGIS Pro.
3. Registering a dataset as branch versioned requires that editor tracking be enabled.
a. True
4. Editing branch-versioned data requires access to a published feature service rather than a direct connection to the enterprise geodatabase.
a. True

Answers to lesson 6 review questions

1. Which type of service may be a candidate for using the shared instance pool rather than dedicated instances?
a. A service for which the minimum dedicated instances is set to zero
2. Users have reported waiting a long time to access a particular service. Which service property must be increased to resolve the issue?
b. Maximum number of instances per machine

Appendix B

Answers to lesson review questions (continued)

3. Which server statistic can be tracked using reports in ArcGIS Server Manager?
 - b. **Maximum running instances**

Answers to lesson 7 review questions

1. When publishing a 3D mesh or a point cloud layer, which ArcGIS Data Store type stores the data?
 - c. **Tile cache data store**
2. A relational data store and a tile cache data store can coexist on a single machine, but the spatiotemporal big data store should be installed on its own machine.
 - a. **True**
3. All hosted services use ArcGIS-managed data, and all services that use ArcGIS-managed data are hosted.
 - b. **False**
(Some hosted services, such as a vector tile layer/service, will use user-managed data. Some non-hosted services, such as map image layer/map services and imagery layer/image services, will use ArcGIS-managed data.)

Answers to lesson 8 review questions

1. When performing analysis in ArcGIS Enterprise, which type of service is always created?
 - b. **Hosted service**
2. How does the object store optimize hosted feature services?
 - c. **Caches query responses**

Appendix B

Answers to lesson review questions (continued)

3. How does the role of distributed data impact scalable distributed analysis?

Due to how the data is stored in multiple parts, the analysis can also be split up into multiple processes. Each part is analyzed using a separate instance of the processing service, which means that one job is being broken up into parts and executed simultaneously, speeding up the time that it takes to execute.

Answers to lesson 9 review questions

1. The spatiotemporal big data store is a requirement to use ArcGIS GeoEvent Server.

b. False

(GeoAnalytics Server must have a spatiotemporal big data store, but GeoEvent Server does not require it.)

2. With which federated ArcGIS Server site is the spatiotemporal big data store registered?

c. Hosting server

3. ArcGIS GeoAnalytics Server can both read and write to a spatiotemporal big data store.

a. True

4. Why might you want to deploy the spatiotemporal big data store across multiple nodes?

Adding machines (nodes) can add data storage capacity, ensure data availability, and allow higher data write throughput.

Answers to lesson 10 review questions

1. A GeoEvent Service can have more than one output.

a. True

2. Sending an email notification when a vehicle leaves a certain geographic area is a task performed by which component of a GeoEvent Service?

b. Output connector

Appendix B

Answers to lesson review questions (continued)

3. Which element of a GeoEvent Service is used to create buffers around GeoEvents after they are ingested into GeoEvent Server?

d. Processor

Answers to lesson 11 review questions

1. A city (as the host) wants to create a collaboration workspace to get data updates from a contractor (as the guest), but it does not need to send data to the contractor. Which access mode should the city choose when creating the workspace?

b. Send only

2. A GIS administrator has an existing distributed collaboration set up between the Parks and Recreation department and the Utilities department for an existing project. The two entities have begun a new project that also requires sharing information between portals. Which action should the GIS administrator perform to support collaboration workflows for both projects?

c. Create a new workspace

3. Each participant in a collaboration links a group from their organization to the workspace.

a. True

Answers to lesson 12 review questions

1. A GIS administrator is working with their IT department to efficiently and effectively deploy ArcGIS Enterprise multiple times across the organization to allow different departments to each have their own ArcGIS Enterprise organization. Which automation strategy should the GIS administrator consider to install repeatable ArcGIS Enterprise deployments?

b. Configuration management

Appendix B

Answers to lesson review questions (continued)

2. Webhooks require ArcGIS Notebooks for automation workflows in ArcGIS Enterprise.

b. False

(Webhooks can be configured in ArcGIS Enterprise portal without a federated ArcGIS Notebook Server licensing role; however, the ArcGIS Notebook Server provides the ability to configure ArcGIS Notebooks as a receiver of a webhook.)

3. The default Publisher role includes the privileges required for portal members to author and run notebooks.

b. False

(Portal members must be assigned a role that has been granted the Create And Edit Notebooks privilege, at minimum. The default Publisher role does not include this privilege.)

4. Which ArcGIS Notebook Server component allows each author to use a subset of your server resources without the author's work or resource use interfering with others?

b. Containers

Appendix C

Additional resources

| Lesson 1 | Resources |
|---|--|
| Architecting an ArcGIS system | <ul style="list-style-type: none">• <i>ArcGIS Architecture Center Help</i> |
| Providing capabilities to support system patterns | <ul style="list-style-type: none">• <i>ArcGIS Enterprise Help: ArcGIS Server 11.3 system requirements</i>• <i>ArcGIS Enterprise Help: ArcGIS Server licensing roles</i> |
| Lesson review | <ul style="list-style-type: none">• <i>ArcGIS Enterprise Help: ArcGIS common patterns of use</i> |

| Lesson 2 | Resources |
|------------------------------------|--|
| Federating additional server sites | <ul style="list-style-type: none">• <i>ArcGIS Enterprise Help: Key concepts for image services</i>• <i>ArcGIS Enterprise Help: Configure and deploy raster analytics</i>• <i>ArcGIS Architecture Center Help: Using system patterns > Working with multiple system patterns (system composition approach)</i> |

Appendix C

Additional resources (continued)

| Lesson 3 | Resources |
|--|---|
| Making ArcGIS Enterprise components highly available | <ul style="list-style-type: none">• ArcGIS Enterprise Help: <i>Configure highly available ArcGIS Enterprise</i> |
| Observing a complex ArcGIS system | <ul style="list-style-type: none">• ArcGIS Monitor - Esri Documentation: <i>Introduction to ArcGIS Monitor</i> |

| Lesson 4 | Resources |
|---|--|
| Explore user-managed data and ArcGIS-managed data | <ul style="list-style-type: none">• Esri technical paper: <i>Data in ArcGIS: User Managed and ArcGIS Managed</i> |

| Lesson 5 | Resources |
|---|---|
| Referencing registered data from ArcGIS Pro | <ul style="list-style-type: none">• ArcGIS Enterprise Help: <i>Make your data accessible to ArcGIS Server</i> |
| Versioning | <ul style="list-style-type: none">• ArcGIS Pro Help: <i>Branch version scenarios</i> |

Appendix C

Additional resources (continued)

| Lesson 6 | Resources |
|---|--|
| Reviewing ArcGIS Server logs and statistics reports | <ul style="list-style-type: none">• ArcGIS Enterprise Help: <i>About specifying server log settings</i>• ArcGIS Enterprise Help: <i>Work with server logs</i>• ArcGIS Enterprise Help: <i>Log codes overview</i> |
| Troubleshooting common service performance issues | <ul style="list-style-type: none">• ArcGIS Server Help: <i>Tune services using best practices</i> |

| Lesson 7 | Resources |
|--|--|
| Managing ArcGIS Data Store using command utilities | <ul style="list-style-type: none">• ArcGIS Enterprise Help: <i>ArcGIS Data Store command utility reference</i> |

| Lesson 9 | Resources |
|---|--|
| Understanding the spatiotemporal big data store | <ul style="list-style-type: none">• ArcGIS Enterprise Help: <i>A quick tour of GeoEvent Server</i> |

| Lesson 11 | Resources |
|---------------------------|--|
| Distributed collaboration | <ul style="list-style-type: none">• ArcGIS Enterprise Help: <i>About distributed collaboration</i> |

Appendix C

Additional resources (continued)

| Lesson 12 | Resources |
|--|---|
| Use ArcGIS Notebooks to schedule administrative tasks in ArcGIS Enterprise | <ul style="list-style-type: none"><li data-bbox="591 449 1422 527">• ArcGIS Enterprise Help: <i>Manage ArcGIS Notebook Server resources</i> |

Contact Esri Training Services

T 909 793 2853 x1585

F 909 793 5953

esri.com/training

380 New York Street

Redlands, California 92373-8100 USA



esri[®]

THE
SCIENCE
OF
WHERE[®]