

ArcGIS Enterprise: Architecting Your Deployment

Assumptions and Prerequisites

This document assumes you are a system architect or an IT administrator (or work as one) and that you are responsible for architecting and/or installing the ArcGIS Enterprise software for your organization.

Before You Begin

Before architecting an ArcGIS Enterprise deployment, you should have a basic understanding of what ArcGIS Enterprise is—the software components that comprise the product as well as the overall capabilities of the software.

Framing Your Existing Product Knowledge

All the software components that make up ArcGIS Enterprise existed in previous generations of the software. Any familiarity you have in working with these components (installing, configuring, etc.) should help you understand how to architect ArcGIS Enterprise. However, particularly if you are moving from an older (end-of-life or near end-of-life) version of the software, be aware that some best practices and recommended architectural patterns have changed. It is important that you always refer to the latest version of the help documentation for updates, deprecations, and other changes in how to get started with the software.

Introduction

There are many things to consider when preparing to deploy ArcGIS Enterprise. A successful and efficient ArcGIS Enterprise deployment has an architecture that has been designed with considerations for

- The capabilities your organization requires.
- How you anticipate your organization will utilize ArcGIS Enterprise.
- The number and type of users you expect for your deployment.
- Clear expectations around service-level agreement (SLA) requirements.

These considerations should be reviewed and revisited at each iteration of planning your deployment, from initiation to scaling for growth. Though it may seem like a daunting task and a heavy responsibility, **don't panic**.

This document is designed to guide your decision-making and expand your product knowledge so that you can design and build the best possible ArcGIS Enterprise deployment for your organization.

Understanding the GIS Spectrum

Before diving into the architecture, it is important that you understand where ArcGIS Enterprise falls within the spectrum of ArcGIS.


While ArcGIS Enterprise is built on the foundation of Server GIS, it is more closely aligned to Web GIS because it enables GIS access through any device, anywhere, anytime—which is the defining characteristic of the Web GIS pattern. As the product continues to evolve, you will likely hear more about distributed GIS—the concept of connecting multiple GIS implementations and sharing beyond a single GIS. But while ArcGIS Enterprise enables distributed GIS, it is not, in itself, a distributed GIS.

- **Desktop GIS**—ArcGIS Desktop (ArcMap or ArcGIS Pro)
- **Server GIS**—A capability of ArcGIS Enterprise powered by ArcGIS Server
- **Web GIS**—ArcGIS Enterprise or ArcGIS Online, working separately
- **Distributed GIS**—ArcGIS Enterprise and ArcGIS Online (or multiple instances of each), working together



Software Components of ArcGIS Enterprise

ArcGIS Enterprise is composed of four software components:



ArcGIS Web Adaptor

This is an Esri-built software load balancer that appropriately directs network traffic and serves as a reverse proxy for Web GIS access; this is also the component that enables web-tier authentication.

Portal for ArcGIS

This is both a web front end and infrastructure back end that supports a user's interaction and overall experience with the Enterprise portal.

ArcGIS Server

ArcGIS Server gives you the ability to publish services and share maps and layers from your own business databases. When used in the base ArcGIS Enterprise deployment, ArcGIS Server is configured as a hosting server that provides the layers and services that power your Web GIS.

ArcGIS Data Store

This is an ArcGIS managed data repository that stores the spatial content that has been copied to the ArcGIS Enterprise portal. It also stores the output of standard analysis tools run within the portal. There are three types of data stores within the ArcGIS Data Store component: relational, tile cache, and spatiotemporal.

ArcGIS Data Store (relational)

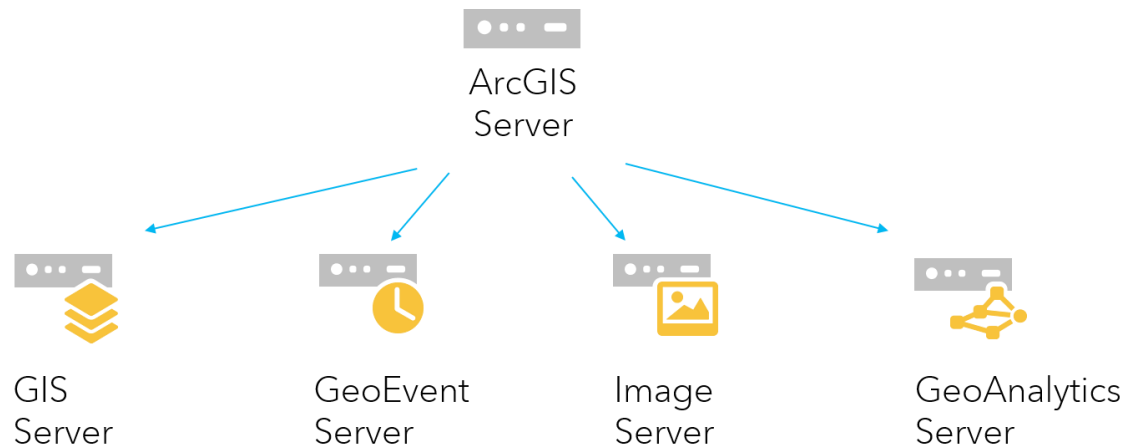
The relational data store is not a replacement for or in competition with enterprise geodatabases that you have configured and administer. RDBMs such as SQL Server, Oracle, or PostgreSQL. Enterprise geodatabases provide a level of control and wealth of functionality that is not exposed with ArcGIS Data Store.

Roles of ArcGIS Server

The ArcGIS Server software component provides different capabilities, depending on how it has been licensed. There are currently five capability-based server licensing roles for ArcGIS Server:

- **GIS Server**
- **Image Server**
- **GeoEvent Server**
- **GeoAnalytics Server**
- **Business Analyst Server**

This document will focus on architectural considerations for GIS Server, Image Server, GeoEvent Server, and GeoAnalytics Server. For architectural best practices related to Business Analyst Server, refer to the Business Analyst Server Help documentation.



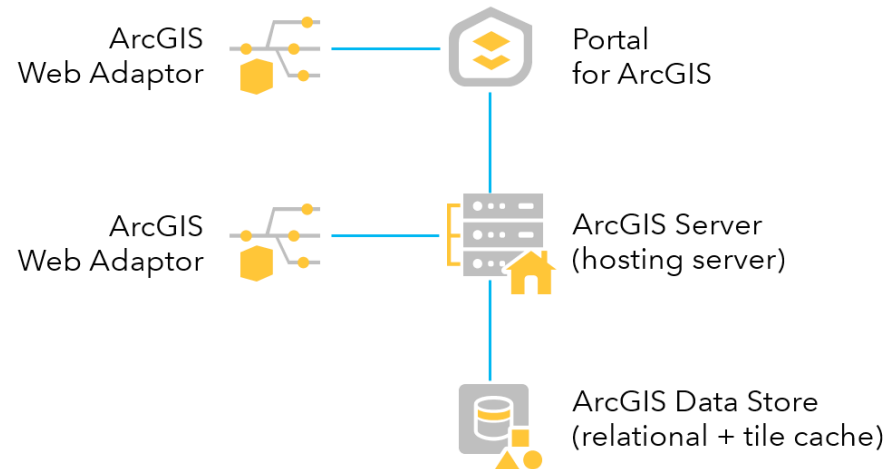
One Software Component, Multiple Server Roles

The Base ArcGIS Enterprise Deployment

The base ArcGIS Enterprise deployment—or, simply, base deployment—is the minimum software configuration for ArcGIS Enterprise. The base deployment is characterized by

- Two ArcGIS Web Adaptors: one for the ArcGIS Server software component and one for the Portal for ArcGIS software component.
- An ArcGIS Data Store instance that has been configured with the relational and tile cache data store types. ArcGIS Data Store is registered with ArcGIS Server as the managed data store.
- ArcGIS Server, which has been configured with the GIS Server licensing role, federated with Portal for ArcGIS, and designated as the hosting server.

Logical Architecture of the Base Deployment



[Base Deployment Configuration](#)

The base deployment can be configured in one of two ways:

- With all the components on a single machine as an all-in-one deployment
- With the components installed on multiple machines as a multi-tier deployment

Functionality of the Base Deployment

A GIS Server powers the base deployment; therefore, all the capabilities you get from ArcGIS GIS Server are accessible with the base deployment. This includes the ability to publish services (such as map services, feature services, and geoprocessing services) with content from referenced data sources including enterprise geodatabases, file-based data (file geodatabases, shapefiles, etc.), and many other spatially enabled databases.

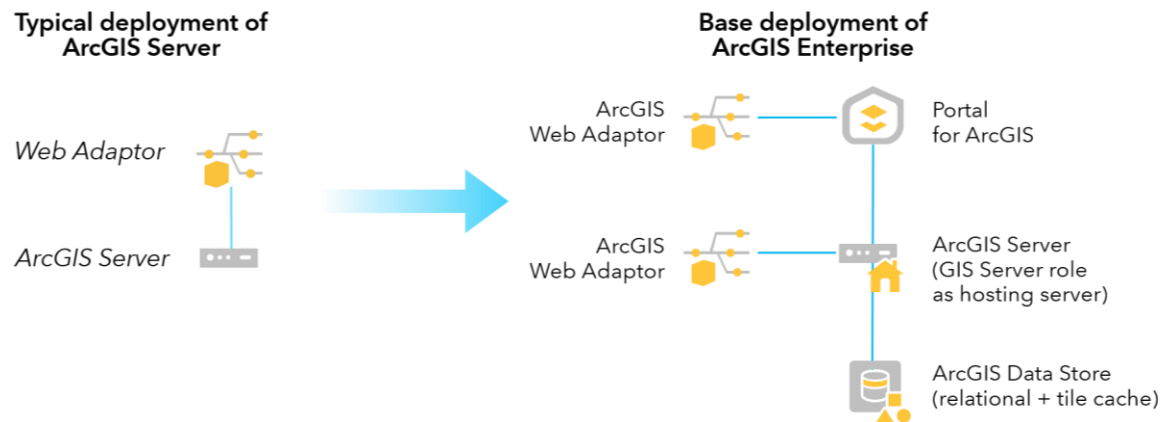
The base deployment is also what brings core Web GIS functionality into your infrastructure, which is why you may have heard ArcGIS Enterprise (or, in previous releases, the Portal for ArcGIS software component) described as "ArcGIS Online in your infrastructure". The core Web GIS functionality, however, is only one part of the overall ArcGIS Enterprise product.

Core Web GIS functionality includes the ability to create apps such as story maps and configurable apps (through app templates or Web AppBuilder for ArcGIS). Using the ArcGIS Enterprise portal as a geospatial content management system also enables self-service mapping, where users can search existing portal content or publish hosted content without requiring direct access to enterprise geodatabases or trusted file shares.

Finally, you can leverage the base deployment in custom apps that you build from the ArcGIS APIs for developers, such as ArcGIS API for JavaScript and ArcGIS Runtime SDKs.

Moving from ArcGIS Server to ArcGIS Enterprise

Though all the software components of the base deployment existed in the software prior to ArcGIS 10.5, many organizations have only deployed ArcGIS Server. If this is the case for your organization, you will need to upgrade your ArcGIS Server and ArcGIS Web Adaptor before adding the remaining components required for the base ArcGIS Enterprise deployment.



Base Deployment Configuration–Quick Steps

1. Acquire SSL certificates and configure HTTPS for your web server.
2. Install, license, and configure ArcGIS Server.
3. Install, license, and configure Portal for ArcGIS.
4. Configure the web adaptors.
5. Install and configure the data store and register it with the server.
6. Federate the server with the portal and designate it as the hosting server.

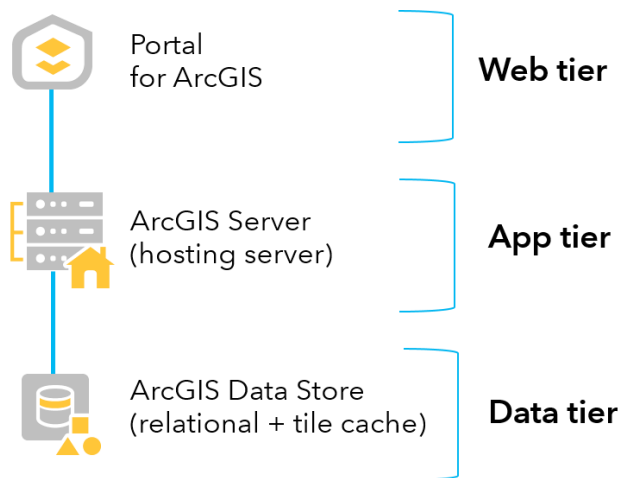
Looking to upgrade your existing ArcGIS Server deployment and aren't quite ready to move to ArcGIS Enterprise? The rest of this document covers the ArcGIS Enterprise deployment architecture, but the [Upgrade ArcGIS Server](#) help topic can assist in upgrading the individual components of your current setup.

For detailed instructions on setting up a base deployment, see the ArcGIS Enterprise Help documentation [Tutorial: Set up a base ArcGIS Enterprise deployment](#).

Choosing a Pattern for Your Base Deployment

The base ArcGIS Enterprise deployment can be configured with all the components on a single machine in what's known as a *single-machine* or *all-in-one* deployment pattern. Alternatively, you can separate the components onto different machines in a *multi-tier* deployment pattern. These two patterns provide the same functionality; however, the multi-tier deployment is better suited to supporting ArcGIS Enterprise deployments that are larger or distributed.

There are three tiers to consider when deciding which base deployment pattern will be the best foundation for ArcGIS Enterprise in your organization: the web tier (Portal for ArcGIS), the app tier (ArcGIS Server), and the data tier (ArcGIS Data Store). Each of these tiers roughly aligns to a component of ArcGIS Enterprise. For each tier, you should think about anticipated usage, load, SLA expectations, and any policy or requirement that would influence the number of machines you should use in your base deployment.



See also [Architecting the ArcGIS Platform: Best Practices](#) for guidance on workload separation and other general architecture advice.

Scaling the Base Deployment

The Enterprise Portal

You rarely need to scale out the web tier of your base deployment (the Portal for ArcGIS software component). Instead of adding new machines to handle the anticipated load on the portal, it is recommended that you simply add resources to the existing machine(s). Using multiple machines in the web tier is typically reserved for setting up a highly available deployment, in which case two machines (a primary machine and a secondary one to prevent a single point of failure) are advised.

Whether or not you create a highly available deployment, you should monitor CPU and memory usage on machines in your ArcGIS Enterprise web tier to see if/when you need to add resources.



ArcGIS Server as the Hosting Server

The hosting server is simply a GIS Server that has been designated as the hosting server through the federation process of configuring the base deployment. Your hosting server site—whether it is made up of one or many machines—is what powers your base deployment.

When it comes time to scale the hosting server site, you have two choices: add more resources to existing machines or add more machines to the site.



Some organizations choose to use the hosting server site, as both hosting server and GIS Server.

There is nothing wrong with this practice. If you choose to use your hosting server as a double-duty server, monitoring CPU and memory usage becomes even more important, and you may need to scale your server site sooner. As a best practice, it is advisable that you implement workload separation and maintain one server site to serve as the hosting server and set up one or more new server sites to handle other GIS workflows. Factors that may cause you to need to scale your hosting server site (even if you are using workload separation) include the following:

- Your users are making heavy use of the built-in analysis tools within the Enterprise portal or tools that are writing to ArcGIS Enterprise from ArcGIS Pro.
- You have a high number of Insights for ArcGIS users, as Insights for ArcGIS will leverage your base ArcGIS Enterprise deployment.

ArcGIS Data Store in the Base Deployment

There are three data store types available from the ArcGIS Data Store component:

Relational—Used for storing hosted data and output from analysis tools

Tile Cache—Used for storing 3D tile caches for 3D visualization

Spatiotemporal Big Data Store—Used for storing large volumes of events from GeoEvent Server and results from GeoAnalytics Server

Only two of these, the relational and tile cache data stores, are used in the base deployment.

The relational and tile cache data stores can be scaled by adding resources to existing data store machines (scaling vertically). Like the other software components of the base deployment, it is important to monitor the ArcGIS Data Store infrastructure—in particular, CPU, memory, disk space, and disk I/O. To better track system performance and monitor for bottlenecks, consider using Esri Professional Services' System Monitor. In late 2017, a fully supported version of System Monitor will be released under the name ArcGIS Monitor- stay tuned!

Deployment scenarios that may require data store scaling include the following:

- Many users requesting data concurrently from many different hosted services—Adding memory is often helpful if this is the bottleneck you are experiencing.
- A large number of concurrent requests—This scenario may cause a CPU bottleneck.

The spatiotemporal big data store is not used in the base deployment. The spatiotemporal big data store can contain large-volume tabular and vector data. It is required to use GeoAnalytics Server and is optional for using GeoEvent Server. More on the spatiotemporal big data store will be covered in later sections.

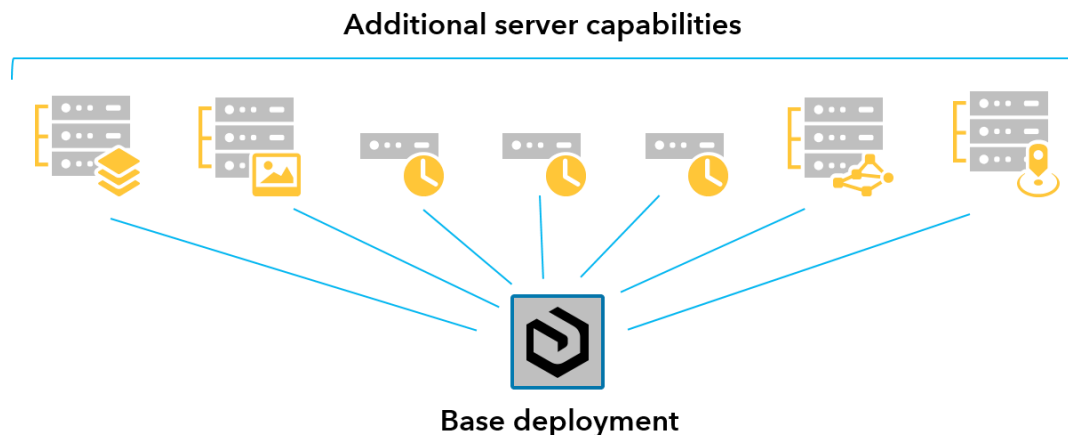
Versions of ArcGIS Data Store prior to 10.5.1 did not always handle out-of-disk-space conditions gracefully. If using a pre-10.5.1 version of the software, be sure to carefully monitor the disk space on your relational data store.

- Use of Insights for ArcGIS—This will consume data tier resources for many operations in order to boost performance.

Expanding beyond the Base Deployment

There are multiple reasons and ways to expand beyond the base deployment. Adding GIS Server sites can provide workload separation by isolating the hosting server from traditional GIS Server duties. When considering workload separation, equally important to isolating the hosting server is separating out the GIS Server instances to support specific high-demand functionality or services—for example, mapping/visualization and heavyweight services such as geoprocessing services.

Besides adding GIS Server sites, you can expand the base deployment by adding new server capabilities through the introduction of a new server role, such as Image Server, GeoEvent Server, or GeoAnalytics Server.



Good to Know

- You can have any number of federated ArcGIS Server sites within your ArcGIS Enterprise deployment.
- Different server roles have different recommendations and restrictions.

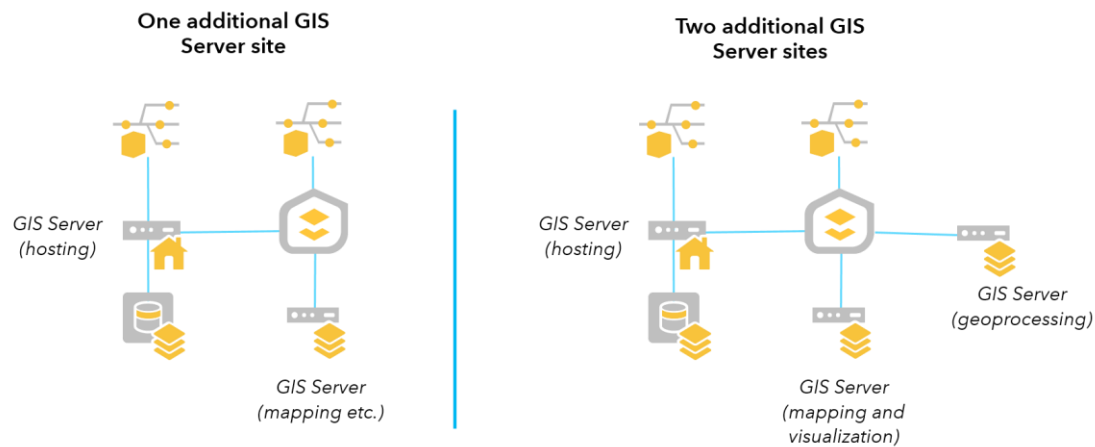
Adding GIS Server Sites to Your Deployment

You already have a GIS Server site as part of the base deployment—it is acting as your hosting server site, which can be made up of one or more machines. However, depending on your deployment and usage scenario, you may need additional sites to achieve optimal performance. You can add as many sites as make sense for your deployment, per the workload separation recommendations.

Types of workload/services that typically benefit from having a separate site include the following:

- Highly used dynamic map services
- Heavyweight geoprocessing services
- CPU-intensive routing services
- Mission-critical services that have different SLAs than other services

Following are sample logical architectures for an ArcGIS Enterprise deployment with multiple GIS Server sites:



Adding Image Server to Your Deployment

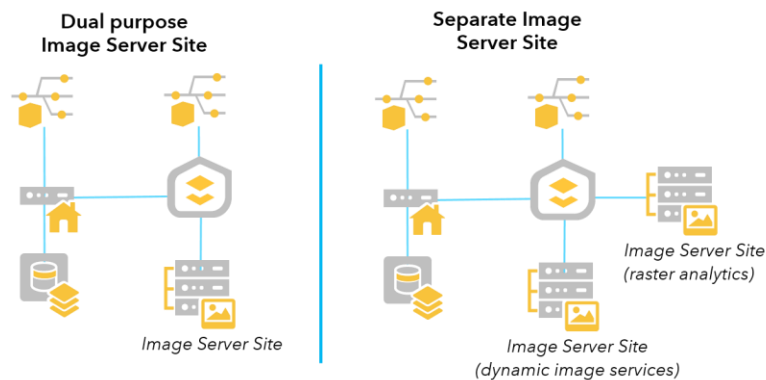
ArcGIS Image Server provides two distinct capabilities:

- Dynamic image services from your own mosaic datasets
- Raster Analytics

The dynamic image service capability allows you to serve large collections of imagery and rasters with dynamic mosaicking and on-the-fly image processing. When this capability is used for dynamic image services, you can have any number of Image Server sites.

Raster Analytics provides the ability to extract information from imagery through distributed processing and analysis of image and raster datasets to create new persisted output at full source resolution. You can only designate one server site in your ArcGIS Enterprise deployment for raster analytics (you do this within the Enterprise portal on the Edit Settings page on the Servers tab).

You can choose to have a separate site for Raster Analytics or use one site for both dynamic image services and Raster Analytics.



Good to Know

At 10.5, the results of Raster Analytics will always be hosted from the site that has been designated for Raster Analytics

This will change in future releases!

Adding GeoAnalytics Server to Your Deployment

ArcGIS GeoAnalytics Server provides distributed computing infrastructure for processing large volumes of vector and tabular data. Within your ArcGIS Enterprise deployment, you can have only one site designated as the GeoAnalytics Server site. As with Raster Analytics, you designate your GeoAnalytics Server site on the Edit Settings page of the Enterprise portal on the Servers tab.

Using GeoAnalytics Server requires the spatiotemporal big data store to be configured with the base deployment. It is important to note that all ArcGIS Data Store types (relational, tile cache, and spatiotemporal) are always registered with the *hosting server site* of the base deployment.

When thinking about how to size/scale GeoAnalytics Server, there are many variables to consider that may impact system performance:


- **The type of data**—File-based vs services based versus big data store.
- **Location of the data**—Local or network.
- **Characteristics of your LAN**—This affects the latency between machines and data store.
- **Characteristics of the machine's data storage system**—This affects disk I/O.
- **Size of Data**—Megabytes vs gigabytes vs terabytes.
- **Ratio of CPU cores to amount of memory per core.**
- **Ability to write result data as fast as it is being generated.**
- Remember to scale the spatiotemporal big data store as needed.

Due to these variables, more machines, more cores, or more memory may not always mean a faster processing time.

Good to Know

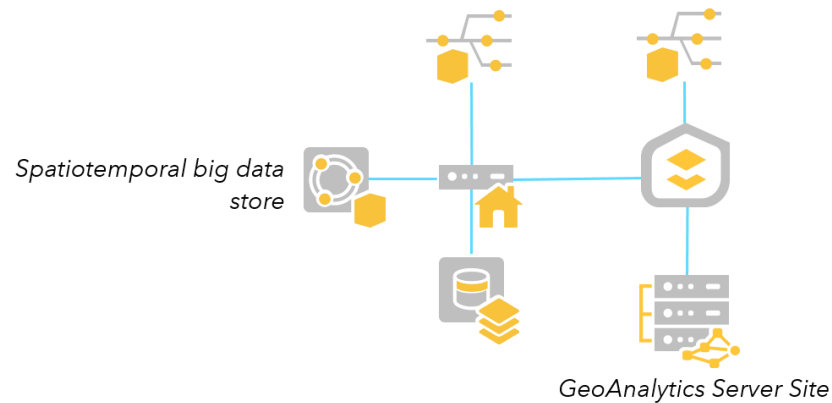
When planning for the spatiotemporal big data store, you will want a machine with at least 16 GB but no more than 64 GB of memory. Memory exceeding 64 GB will not be utilized by the software.

Below is a sample logical architecture for adding GeoAnalytics Server to your base deployment:



At minimum, you will need three machines to support an ArcGIS Enterprise deployment with GeoAnalytics Server: one for the base deployment, one for the GeoAnalytics Server site, and one for the spatiotemporal big data store.

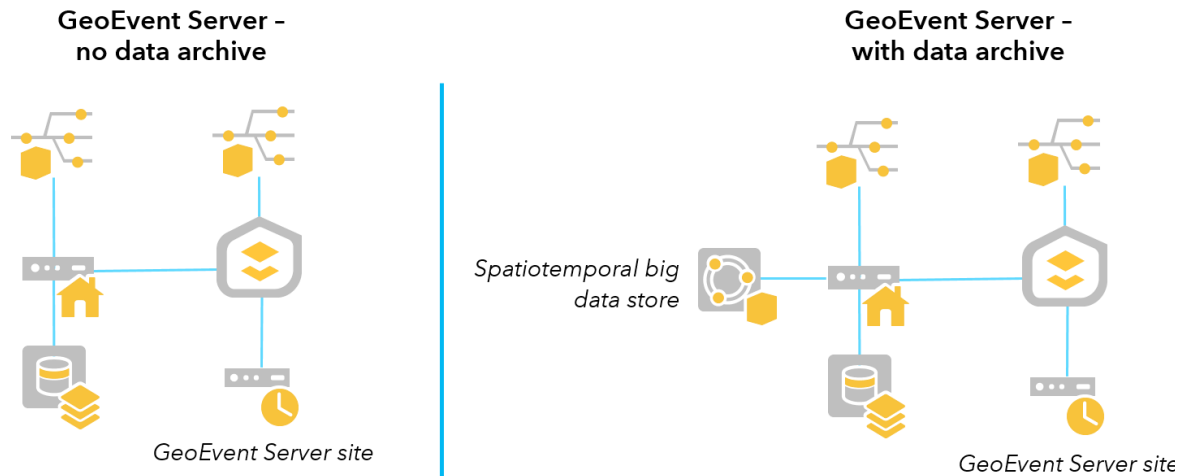
If/When you are ready to scale your GeoAnalytics Server site, rather than add another site, simply add machines/resources to your existing GeoAnalytics Server site:



Adding GeoEvent Server to Your Deployment


ArcGIS GeoEvent Server provides the ability to configure event processing workflows, “GeoEvent services”, which apply real-time analytics and filtering to ingested event records. At ArcGIS Enterprise 10.5.1 and prior versions, it is strongly recommended that you use only single machine sites for GeoEvent Server. Each machine that runs GeoEvent Server must be powerful enough to handle peak throughput for the total number of event records ingested. Depending on the number of event records you anticipate ingesting and processing each second, you may need to scale up your GeoEvent Server site.

If you would like to archive event data processed by GeoEvent Server, use the spatiotemporal big data store.



To scale GeoEvent Server, do not add machines to an existing GeoEvent Server site. Instead, you should use additional, separate GeoEvent Server sites. Using multiple single-machine GeoEvent Server sites will allow you to scale your solution horizontally, sending different data streams to instances of GeoEvent Server dedicated to handling event records from those streams.

See the following example of logical architecture for a scaled out GeoEvent Server deployment with the spatiotemporal big data store for data archiving:



Recap: Number of Server Sites per Deployment

GIS Server

- As many sites as make sense for your deployment, following workload separation recommendations

Image Server

- As many sites for dynamic image services as make sense
- One site for Raster Analytics

GeoAnalytics Server

- Only one site per deployment

GeoEvent Server

- As many sites as make sense (If using ArcGIS 10.5.x or prior version, scale by adding multiple single-machine sites.)

Deploying ArcGIS Enterprise—Overview

Deploying ArcGIS Enterprise does not need to be challenging, nor does it need to be a manual process. A suite of free-to-use tools will help you jump-start your deployment.

ArcGIS Enterprise Builder

ArcGIS Enterprise Builder is a wizard that installs and configures a base ArcGIS Enterprise deployment on a single machine. ArcGIS Enterprise Builder is available for download from My Esri.

Cloud Tools

If you are deploying in the Amazon Web Services (AWS) or Microsoft Azure clouds, Esri provides machine images and builders that will install and configure the ArcGIS Enterprise software and spin up the underlying infrastructure of your choice. These tools can be found on each cloud provider's respective marketplace and follow a bring-your-own license model.

Chef

Chef is an open-source scripting language for automation. Chef is a great choice for automating the installation and configuration of larger and more advanced (multi-tier) deployments. The ArcGIS Enterprise development team has created several Chef cookbooks and recipes available on the Esri GitHub repository to help you get started.

Updating Best Practices

It is important to always review the product documentation, even if you are/were familiar with the previous release of a product. Several key updates, which are contrary to previous practices, have been made to best practices and recommendations for architecting the software.

For example:

Do not use unfederated siloed ArcGIS Server sites.

- Esri has made significant improvements in the performance and stability of multimachine sites since ArcGIS 10.3.1; also, siloed sites do not work in the base deployment (the federated model) or with ArcGIS Pro.

Do not use clusters.

- The use of clusters on any version of the software is discouraged. In fact, support for clusters has been officially deprecated as of 10.5.1 and will be removed from the software entirely in the next several releases.

Do federate your ArcGIS Server sites.

- Federation between the ArcGIS Server and Portal for ArcGIS software components is required in order to set up the base deployment.

Do not unfederate your ArcGIS Server sites when upgrading.

- Unfederating ArcGIS Server sites is an irreversible process that is not needed during ordinary workflows. Follow the upgrade workflow outline in the documentation.

Do use ArcGIS Web Adaptor.

- The ArcGIS Web Adaptor simplifies a lot of configuration that would otherwise have to be done manually. Unless you have a good technical reason not to, always use the web adaptor (one for ArcGIS Server and one for Portal for ArcGIS).

Dispelling Old Myths

Just as best practices are updated, there are certain beliefs about software operating systems that are no longer true:

Myth: "Windows can only run approximately 250 service instances."

(referring to SOC processes)

Truth: This is a Windows imposed limitation that can be lifted by changing a setting within the Windows registry. Remember! You must still have enough system resources (memory and CPU) to power all these services. Unless you have at least 50 GB RAM available for running services, you are unlikely to be impacted by this.

See this technical article for more information and specific steps:

support.esri.com/technical-article/000001218

Myth: "Linux can only run 250 service instances." (referring to SOC processes)

- **Truth:** ArcGIS Server 10.3.x had such a limitation. If you upgrade to ArcGIS 10.4 or higher, you will not experience this limitation.

Key Takeaways

- Stay up-to-date on best practices—they change over time!
- ArcGIS Enterprise is designed using the federated server model.
 - Features that require federation include the following:
 - Publishing from ArcGIS Pro
 - Raster Analytics (Image Server)
 - GeoAnalytics Server
 - Using the spatiotemporal big data store
- Understand the base deployment.
- Understand the individual server roles and the recommendations and requirements of each—they're not all the same!
- Remember, **don't panic**.