

Reference Architecture for a Parcel Management System: Local Government

Last generated: July 10, 2025



Table of Contents

- **Reference Architecture**
 - Introduction and overview 2
 - Intended workflows 4
 - Logical architecture 6
 - Design choices and considerations 9
 - Using this reference architecture 14

Parcel Management System

This reference architecture represents a foundational Parcel Management System for a local or state government, or any other organization that manages parcel and cadastral data. It is based upon the [data management and editing system pattern](#) and designed for deployment with ArcGIS Enterprise on Windows or Linux. It includes support for management of 3D and 4D parcel data, including strata and subsurface information, ownership records, and agricultural and natural resource rights, using the [ArcGIS Parcel Fabric](#).

- Learn [about reference architectures](#) first to gain important context.

Note:

This reference architecture represents an example blueprint for this system type. The design does not reflect the needs of all Parcel Management Systems, including those that may [combine multiple system patterns](#).

Successful system design involves the selection of workflows used to define requirements as well as design both logical and physical architectures. This reference architecture is based on a specific set of workflows identified for Parcel Management Systems, which are described in the [intended workflows](#) section.

Capabilities

A foundational Parcel Management System delivers all capabilities as defined in the [data editing and management system pattern](#). It focuses on delivery of parcel management capabilities that include:

- Service-based editing of structured relational data stored in an enterprise geodatabase, including geometry, attributes, and relationships.
- Data collection workflows using forms and maps in ArcGIS Pro, optimized for productivity.
- Complex parcel fabric management using the ArcGIS Parcel Fabric as an advanced data model.
- 3D and 4D parcel data management, including strata and subsurface information, ownership records, agricultural, natural resource rights and time-based information like ownership changes.
- Interacting with, querying, and viewing the parcel fabric through web applications and experiences.

Software components

This reference architecture has been constructed using recent versions of the following ArcGIS software components. The system also supports connections from and interactions with other Esri software or client applications, but the workflows and testing of this system were focused on these components:

- [ArcGIS Pro](#)
- [ArcGIS Enterprise](#), including [ArcGIS License Manager](#)
- [ArcGIS Monitor](#)
- [ArcGIS Online](#)

The role of each of these components is further detailed in the [logical architecture](#).

Intended workflows

This reference architecture was designed based on a specific set of intended workflows that users of such a system would be completing on a regular basis. A workflow in this case refers to the series of tasks that are performed by a user (or users) of the system to achieve a specific business process or objective.

For example, a workflow performed on a parcel fabric to create a feature would include tasks such as:

- Search for the parent parcel
- Zoom to the location of interest (LOI)
- Create a parcel fabric record
- Create the feature via COGO ([Coordinate Geography](#))
- Build the parcel
- Align the parcel to surrounding features
- Validate the topology
- Reconcile and post data to a feature service

Workflows can vary significantly from deployment to deployment and the workflows used in this reference architecture are intended to represent common versions that are straightforward and general enough to represent many different potential deployments.

Workflows are an integral part of the reference architecture definition as they help define the requirements that the architecture needs to meet. Additionally, clear, and specific workflow definition helps to:

- Define the necessary system functionality and the appropriate applications that end-users will interact with.
- Inform system requirements around architecture pillars like [performance](#), [reliability](#), and [security](#).
- Inform physical system requirements such as CPU, memory, storage, and networking, which directly impact machine types, sizes, and hardware configurations.

To ensure that workflows are useful for system design purposes, they need to represent real user experiences. This reference architecture was designed with a focus on the [editor persona](#) along with use cases for a general user persona. With these audiences identified, the following workflows were selected and developed:

Editor persona

Record-driven workflows

1. Create a feature by entering in new survey points or lines.
2. Conduct a boundary adjustment with the [Align](#), [Split](#), or [Merge](#) tools. Take existing feature(s) and change their legal boundaries, retiring the underlying parcel(s).
3. Attribute adjustment – Modify non-spatial attributes, like a transfer of title.
4. Re-plat – Manage changes to the property or properties where a new survey was created.

Quality-driven workflows

1. Conduct spatial quality analysis, such as running a least squares analysis.

General user

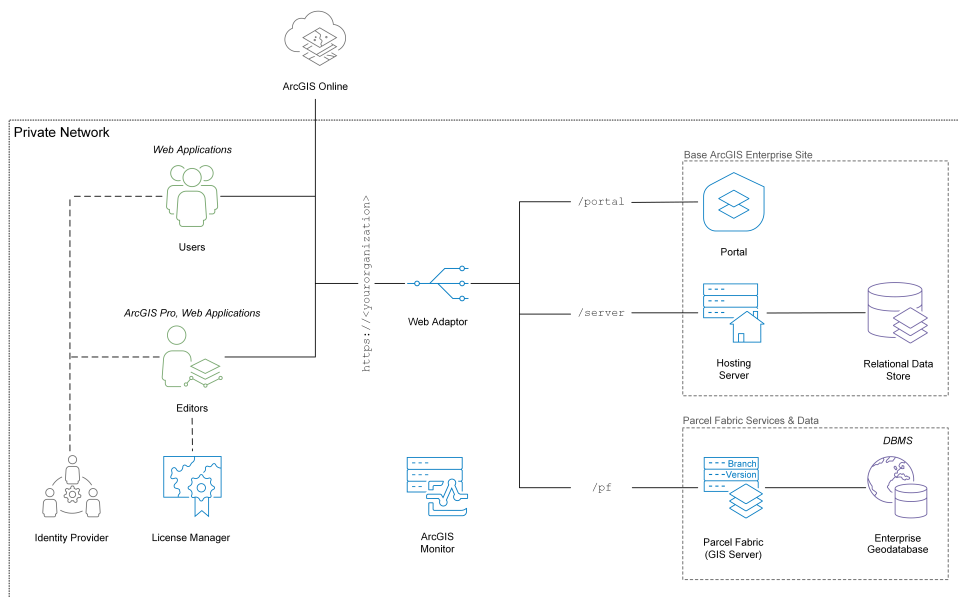
1. Query features - such as searching for parcels and displaying historic chain of title.
2. View and interact with the parcel fabric - as an appraiser would when performing a value analysis within a neighborhood.

Note:

An implementation of this Parcel Management System reference architecture may include other types of workflows not listed here, such as mobile data collection, which would introduce additional architectural considerations. See [related system patterns](#), such as the data editing and management system pattern, for more information.

Logical reference architecture

This logical view of the reference architecture defines individual software components, how they are separated or combined, and key interactions between them.



Download a [Microsoft Visio file](#) of this architecture.

Learn more about [diagramming resources](#) for ArcGIS systems.

Note:

This reference architecture was designed with a focus on specific [workflows](#) and [capabilities](#).

Software components

This diagram represents a variety of software components that contribute to the overall system design, including:

- Apps to provide end-users with access to system capabilities. The primary apps used in Parcel Management Systems are:
 - [ArcGIS Pro](#) - a professional GIS desktop application used by editors
 - Configurable apps and app builders - used by general users and editors
- An ArcGIS Enterprise [base deployment](#), consisting of four components:
 - [Portal for ArcGIS](#)

- [ArcGIS Server](#) configured as a hosting server
- [ArcGIS Data Store](#) configured as a relational store for persisting [ArcGIS-managed](#) data
- [ArcGIS Web Adaptor](#) to provide load balancing and reverse proxy capabilities. Additional load balancing and reverse proxy components may be recommended or required for certain scenarios, including [high availability](#) deployments.
- An additional ArcGIS Server site configured with the GIS Server role and used to deliver branch versioning and parcel fabric editing services.
- An [enterprise geodatabase](#) to store and manage an ArcGIS Parcel Fabric through a [supported relational database management system](#) (DBMS).
- An enterprise identity provider. Though not strictly required, this component is recommended to provide users with a single sign-on (SSO) experience when using ArcGIS Enterprise. Learn more about [ArcGIS authentication models and providers](#).
- [ArcGIS License Manager](#) to [configure and manage ArcGIS Pro licenses](#) used by desktop editors. This component is no longer required for deployments of ArcGIS Enterprise 11.4 or newer.
- [ArcGIS Monitor](#) to monitor and optimize the system components and overall health.
- [ArcGIS Online](#), providing basemaps and other location services.

Key interactions

The software components described above interact with each other in the following ways:

- Client applications communicate with enterprise data services over HTTPS, typically via stateless REST APIs. A friendly, well-defined domain name is recommended as the entry point to the system. Three separate ArcGIS Web Adaptor instances configured in the web server handle context path-based routing to the Portal for ArcGIS and ArcGIS Server components described above. Learn more about [DNS, naming, and URLs](#).
- Client applications communicate with basemaps, and location services provided by ArcGIS Online over HTTPS, typically via stateless REST APIs. This requires connectivity from client machines to the internet.
- ArcGIS Server maintains persistent TCP connections to both the database management system (DBMS) hosting the enterprise geodatabase and ArcGIS Data Store. The former requires that [appropriate database client software](#) be installed on the ArcGIS Server machines communicating with the DBMS.

- ArcGIS Pro communicates with the ArcGIS License Manager using TCP/IP protocols. Learn more about [configuring ArcGIS License Manager with ArcGIS Enterprise Portal](#) and configuring [ArcGIS License Manager to work through a firewall](#).
- ArcGIS Monitor communicates with a variety of ArcGIS and IT components, such as RDBMSs, using a variety of patterns and protocols. See [ArcGIS Monitor documentation](#) for more information.
- Location services hosted and managed by ArcGIS Online are typically registered and made available for use within ArcGIS Enterprise. See [configuring ArcGIS Online utility services](#), [configuring ArcGIS Living Atlas content](#), and [distributed collaboration](#).

Additional information on interactions between ArcGIS Enterprise components can be found in the ArcGIS Enterprise on Windows and Linux [product documentation](#), including a [diagram of ports](#) used in an ArcGIS Enterprise on Windows and Linux deployment.

Design choices and considerations

The design of this reference architecture assumes that the Parcel Management System is a business and/or mission-critical enterprise system, where requirements include:

- Minimal system downtime (with zero data loss) due to expected or unexpected events
- Excellent service performance, without substantial lag that would hinder end-user productivity
- An efficient, effective, and overall positive end-user experience

To achieve these goals, the design needs to meet certain criteria within each [architecture pillar](#). Some of the recommended practices associated with each pillar are outlined below, but they do not represent a complete set of architectural considerations. Please refer to the data editing and management system pattern's [considerations](#) for more information.

Performance and scalability

In terms of [performance and scalability](#), this architecture aims to optimize the overall experience users have with the system, while responding to evolving workload demands. A Parcel Management System should deliver editing experiences with consistent performance metrics, to create a positive end-user experience that increases end-user efficiency. In addition to the performance improvement practices outlined below, relational database performance management is also a major factor in the overall performance of your Parcel Management System.

Workload separation

[Workload separation](#) is a design approach focused on optimal distribution of compute resources. For example, some editing requests in a Parcel Management System may take longer to process than standard map requests, so editing workloads may benefit from separate, dedicated compute resources (such as an ArcGIS GIS Server site). This approach of workload separation helps separate long-running requests from shorter requests so that editors have dedicated resources and viewers are not impacted by long transactions. If this separation is implemented, the system performance for both groups is likely to improve as resource contention is reduced, and the system is more easily scaled. Resources can be added to either server site to scale horizontally or vertically.

Workload separation can be implemented in different ways, such as:

- **By component** - Separating components onto different virtual machines or compute infrastructure ensures that individual components are not contending for system resources. While it is supported to install and configure multiple ArcGIS Enterprise components on a single system, it is generally not recommended in well-architected production systems.
- **By service type** - Separating ArcGIS Server components by service type is another workload separation approach that can be applied to this architecture. For example, using separate GIS Server sites to support ArcGIS Parcel Fabric workloads and hosted service or mapping workloads.

Co-location

Co-location is a design approach where system components are deployed to the same data center, in the same sub-network, which helps reduce network latency by reducing the communication distance across the network. In general, network latency is more impactful to end-user experience than network bandwidth for common GIS operations. Another consideration in this area is the location of user and client machines. If a user has a high latency connection, the co-location of system components is unlikely to improve their experience of working with the system. In some cases, use of thin clients or remote access may be preferable to physical hardware that connects over a slow or overloaded network.

Reliability

[Reliability](#) ensures your system provides the level of service required by the business, as well as your customers and stakeholders. As a business or mission-critical enterprise system, Parcel Management Systems always require backups of the data, and often require backups of system components. Although data correctness is generally more critical than system uptime for Parcel Management Systems, they may also require a high availability configuration to meet business needs.

Backups

For enterprise systems with availability expectations, requirements, or commitments, a clearly defined, actionable, and well-tested backup approach is critical. With Parcel Management Systems, data-level backups of the ArcGIS Parcel Fabric are essential, at a minimum. Depending on an organization's requirements, backups of other system components may also be needed. Refer to [backups and disaster recovery](#) for more information on backup strategies and methods.

High availability

[High availability](#) is a design approach that aims to build the system to meet a prearranged level of operational performance over a specific period. A highly available system needs redundancy, system monitoring, and automation commensurate with the target service level agreement (SLA).

Redundancy might include disparate components such as network connectivity, power reliability, data center cooling, and access to staff with the skills to maintain the system. Automation might be designed to take action based on monitoring to avoid outages. For more information, see

[configuring highly available ArcGIS Enterprise components](#).

Note:

Keep in mind that high availability configurations significantly increase infrastructure and operational costs of the system and requires specialized skills to ensure its success. High availability designs require an operational commitment across people, process, technology, and governance.

Observability

[Observability](#) provides visibility into the system, enabling operations staff and other technical roles to keep the system running in a healthy, steady state. Monitoring of system availability, performance, and usage is critical to a Parcel Management System. In addition to monitoring the ArcGIS Enterprise software, it is important to [monitor all supporting components and infrastructure](#), such as the Windows or Linux operating system, databases and other data stores, compute, network, security perimeter and any other relevant components.

Monitoring

Any organization must have an enterprise IT [monitoring](#) and response framework in order to successfully build and operate enterprise systems. Proactive monitoring of systems is as important as reactive problem-solving, and effective capture of telemetry provides awareness of the system at any given time and identifies trending system behaviors.

ArcGIS Enterprise on Windows/Linux can be observed in a variety of ways including server logs and server statistics. In addition to monitoring the ArcGIS Enterprise software, it is important to monitor all supporting components and infrastructure such as the Windows or Linux operating system, databases and other data stores, as well as compute, network, security, and other infrastructure.

In the reference architecture, the telemetry capture mechanism is represented by ArcGIS Monitor. Some important system characteristics captured in this way include:

- CPU usage
- RAM consumption
- Disk activity
- Network activity
- ArcSOC usage and saturation in GIS Server sites

Capturing Telemetry

A key aspect of observability is the use of [telemetry](#), data or information that represents real user activity on a system. Capturing telemetry across all components of the design, including desktop client machines, is critical to understand the performance and utilization of the system, with the overall intent of identifying bottlenecks and opportunities for optimizing the system.

It's important to note that telemetry outputs do not necessarily capture the full user experience. Consider engaging with end-users of applications (such as those using desktop, web, or mobile apps) to observe their experience performing their workflows in addition to telemetry capture.

Security

[Security](#) protects your systems and information. Security design considerations for a Parcel Management System are closely aligned to the [system pattern security requirements](#), including important considerations for user authentication, system authorization, data and access control, and auditing of user activity and system configuration changes.

Integration

[Integration](#) connects this system with other systems for delivering enterprise services and amplifying organizational productivity. A Parcel Management System typically needs to accommodate data exchange and alignment with other systems like a Computer Aided Mass Appraisal (CAMA) Systems, Recording or Registry systems, and/or other Land Management solutions. Integration requirements for a Parcel Management System are closely aligned to the [data editing and management system pattern](#).

Automation

[Automation](#) aims to reduce effort spent on manual deployment and operational tasks, leading to increased operational efficiency as well as reduction in human-introduced system anomalies.

Automation requirements for a Parcel Management System are closely aligned to the [data editing and management system pattern](#), and include practices like:

- Extensive workflow automation, such as with tasks in ArcGIS Pro.
- Use of Python scripting to automate repeatable data management tasks, like digital submission, map sheet generation, and QA/QC checks.
- System administration automation, including software deployment, use of infrastructure as code (IaC), and a DevOps approach to any custom application development.

Physical design considerations

Esri offers [system architecture design](#) services should you need help determining all of the different factors relating to your organization's physical design, such as networking, storage, system environments, and sizing.

Additional resources related to physical design include:

- [High availability](#)
- [Workload separation](#)
- [Environment isolation](#)
- [Load balancing and reverse proxies](#)
- [Secure network design](#)
- [Storage considerations](#)
- [Virtualization and ArcGIS Enterprise](#)
- Minimum system requirements, listed in [ArcGIS product documentation](#)

Using this reference architecture

This reference architecture has been prepared and presented as a part of the ArcGIS Architecture Center to help organizations understand how they might combine ArcGIS components together in a standardized and well-defined configuration. To make best use of this reference architecture, the following approach is suggested:

1. **Assess** the [workflows](#) used as the focus of the reference architecture to determine their suitability to your own future system.
2. **Use** the [logical reference architecture](#) as a starting point to design your Parcel Management System. Where your workflows or requirements differ, the [design process](#) can provide useful considerations to resolve these differences and introduce new components as needed.
3. **Incorporate** [design choices and considerations](#) with respect to your organization's non-functional requirements in areas like [performance and scalability](#), [reliability](#), and [observability](#).
4. **Learn** more about the [architecture practices](#) illustrated in this reference architecture.

Esri offers [system architecture design](#) services should you need help in the design process.

Additional related resources

These resources reference key concepts and practices that may not have been explicitly stated within the reference architecture, but can be significant contributors to successfully designing, implementing, and operating a Parcel Management System.

- [Build and manage a GIS program](#)
- [IT Governance](#)
- [Upgrades and patching](#)
- [Backups and disaster recovery](#)
- [Secure network design](#)