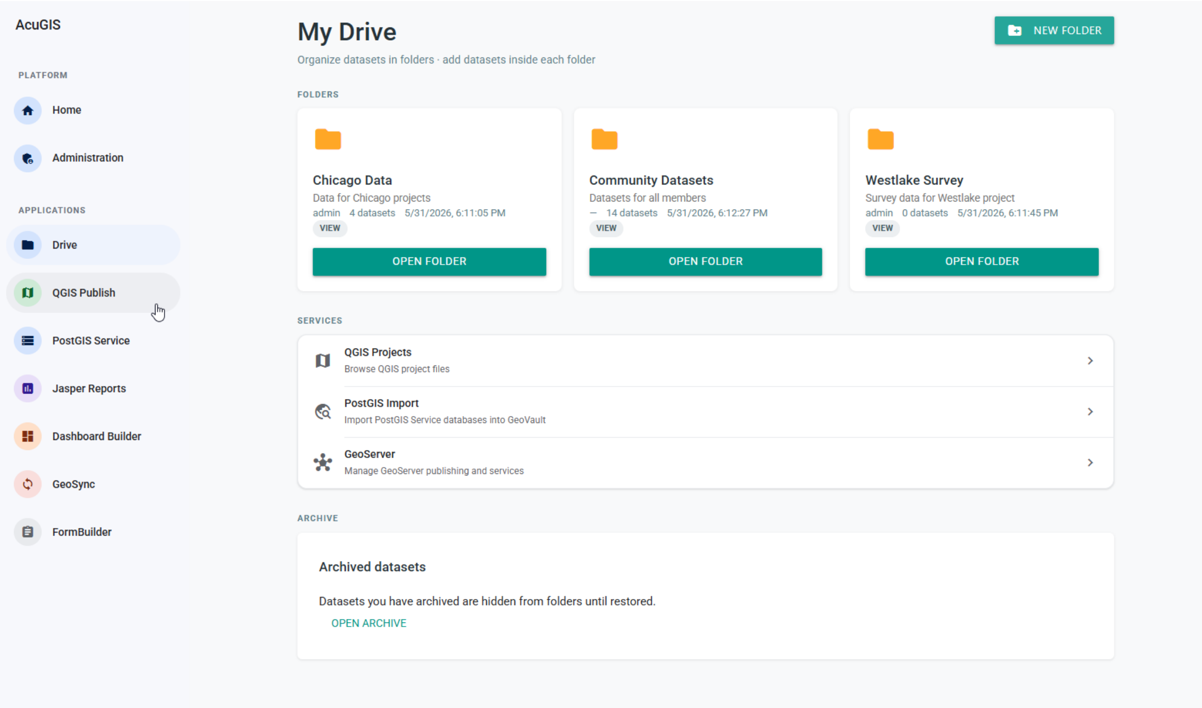

AcuGIS Cloud: Deterministic GIS Pipelines and Spatial Operations

A Unified Self-Hosted Spatial Data Platform

Executive Summary

- Why GIS workflows remain fragmented
- The problem of non-reproducible spatial workflows
- Platform overview
- Deterministic spatial operations as a design principle
- Key capabilities and intended audiences

1. The GIS Workflow Problem

1.1 Fragmented Spatial Systems

1.2 Manual and Non-Repeatable GIS Workflows

1.3 Mobile Data Collection Challenges

1.4 Publishing and Reporting Disconnects

2. Platform Architecture Overview

2.1 System Architecture

High-level architecture diagram

2.2 Core Components

- GeoETL
- GeoSync
- GeoVault
- PostGIS Service
- QGIS Publisher
- GeoServer integration
- TileServer
- Reporting / Jasper

2.3 Design Principles

- Deterministic execution
- Reproducibility
- Self-hosted deployment
- Open standards
- GIS-native architecture
- Human-guided automation

5.4 Editing and Operational Management

5.5 Metadata and Discovery

6. PostGIS Service

6.1 GIS Database Lifecycle

6.2 Provisioning and Access Control

6.3 Schema and Database Management

6.4 Administrative Workflows

7. Publishing Architecture

7.1 QGIS Publisher

7.2 GeoServer Integration

7.3 TileServer

7.4 Web Service Delivery

- WMS
- WFS
- Vector tiles
- Raster services

12.1 Desktop-Centric GIS

12.2 Manual ETL Workflows

12.3 Commercial Spatial ETL Platforms

12.4 Self-Hosted Spatial Operations Platforms

Conclusion

Why deterministic GIS pipelines matter.

Executive Summary

Geographic information systems (GIS) remain foundational to government, utilities, environmental monitoring, infrastructure management, logistics, and enterprise analytics. Yet despite advances in web mapping, cloud infrastructure, and spatial databases, operational GIS workflows remain fragmented across disconnected tools and manual processes.

Organizations commonly rely on separate systems for data acquisition, transformation, field collection, synchronization, storage, publishing, and reporting. Spatial data may originate from files, APIs, cloud storage, enterprise databases, or field collection systems, yet moving data from acquisition to publication frequently requires repeated manual intervention, desktop GIS processing, custom scripts, or brittle ETL chains. These workflows are often difficult to reproduce, difficult to automate, and difficult to audit.

The challenge is not merely one of tooling, but of reproducibility.

Many spatial workflows remain procedural rather than deterministic: transformations are applied manually, field synchronization processes vary between projects, publishing pipelines are disconnected from ingestion systems, and operational reporting depends on bespoke integrations. As datasets grow and organizations increasingly require repeatable, auditable workflows, the limitations of fragmented GIS architectures become more pronounced.

This paper presents a unified, self-hosted spatial operations platform designed to address these challenges through deterministic GIS pipelines and operational integration.

The platform spans the complete lifecycle of spatial operations:

Acquire → Sync → Shape → Store → Publish → Report

Core components include:

- **GeoETL** for deterministic spatial ingestion, inspection, transformation, and scheduled execution pipelines
- **GeoSync** for operational synchronization of mobile and field-collected spatial data into enterprise PostGIS environments
- **GeoVault** for dataset lifecycle management, cataloging, versioning, and operational editing
- **PostGIS Service** for spatial database provisioning and lifecycle management

- **QGIS Publisher**, GeoServer integration, and tile services for web publication and spatial service delivery
- **Reporting and dashboard systems** for scheduled operational reporting and downstream analytics

Rather than emphasizing opaque automation, the platform prioritizes deterministic execution, inspectable workflows, and reproducibility. Natural-language assistance may be used to simplify transformation design, but execution remains explicit, reviewable, and repeatable.

The resulting architecture enables organizations to standardize spatial operations while remaining fully self-hosted, interoperable with existing GIS infrastructure, and aligned with open geospatial standards.

This document describes the architecture, operational model, workflows, and implementation approach of a modern spatial operations platform designed for repeatable GIS workflows at enterprise scale.

1. The GIS Workflow Problem

1.1 Fragmented Spatial Systems

Modern GIS environments are rarely unified systems. Instead, they are composed of independent tools assembled over time to address discrete operational needs.

Spatial data may originate from downloaded shapefiles, GeoJSON, GeoPackages, APIs, cloud object storage, enterprise databases, ESRI Feature Services, or mobile field systems. Data preparation may occur inside desktop GIS software, custom Python scripts, SQL procedures, spreadsheets, or ad hoc command-line utilities. Publishing frequently relies on separate map servers, while reporting and dashboards are implemented independently of data ingestion workflows.

A typical operational workflow may resemble the following:

1. Acquire source data from a file, URL, API, or database
2. Clean or reshape the data in desktop GIS or scripts
3. Import into PostGIS
4. Publish layers to web services
5. Synchronize field-collected updates
6. Generate reports or dashboards

Although each step is individually manageable, the workflow as a whole is often fragile and difficult to reproduce.

Transformations applied manually in desktop environments may not be documented. Data synchronization methods differ between projects. Publishing logic is frequently coupled to server-specific configuration. Reporting systems often consume data only after custom integration work.

The result is operational fragmentation: systems that function independently but fail to operate cohesively.

1.2 Manual and Non-Repeatable GIS Workflows

Many GIS workflows remain highly procedural.

A user may download a dataset, rename fields, repair geometry, reproject coordinates, normalize schemas, filter records, export a revised file, import into PostGIS, configure publishing, and repeat the process later when updated data arrives.

The challenge is not simply the number of steps. The challenge is that the workflow itself is frequently undocumented or non-repeatable.

Two analysts processing the same source data may produce different outputs. Operational knowledge often exists only in institutional memory or undocumented scripts. When source data changes, workflows must be reconstructed manually.

This introduces operational risk:

- Inconsistent outputs
- Limited auditability
- Reduced reproducibility
- Difficult automation
- High onboarding cost for teams
- Increased dependence on individual operators

In enterprise and public-sector GIS environments, repeatability increasingly matters as much as analytical capability.

Organizations require workflows that can be inspected, reproduced, scheduled, validated, and operated consistently across teams and projects.

1.3 Mobile Data Collection Challenges

Field data collection introduces an additional layer of complexity.

Mobile GIS systems enable collection of inspections, environmental observations, infrastructure edits, and operational updates in disconnected environments. However, synchronization of field-collected data into enterprise databases often remains inconsistent or highly project-specific.

Teams may rely on manual exports, intermittent imports, custom synchronization scripts, or ad hoc reconciliation workflows.

Common challenges include:

- Delayed synchronization of field edits
- Manual import/export processes
- Dataset duplication
- Inconsistent update handling
- Difficulty integrating mobile workflows with enterprise databases

The challenge is not simply collecting field data; it is operationalizing field data.

A spatial operations platform must provide predictable, repeatable synchronization between mobile workflows and enterprise spatial systems.

1.4 Publishing and Reporting Disconnects

In many GIS organizations, publication is treated as a separate concern from ingestion.

Data is imported into databases, then separately configured for publication through map servers or desktop publishing workflows. Reporting systems are frequently disconnected from both ingestion and publication layers.

This separation increases operational complexity:

- Published services may diverge from operational datasets
- Reporting systems depend on custom integrations
- Data refresh processes become difficult to automate
- Publishing logic becomes tightly coupled to infrastructure

The result is a fragmented operational lifecycle rather than an integrated spatial workflow.

A modern GIS platform should not merely publish maps. It should provide repeatable operational workflows spanning acquisition, synchronization, transformation, storage, publication, and reporting.

2. Platform Architecture Overview

2.1 AcuGIS Cloud Overview

AcuGIS Cloud is a unified spatial operations platform designed to manage the complete lifecycle of geospatial data, from acquisition and field synchronization through publication and reporting.

Traditional GIS environments often rely on multiple disconnected systems for data ingestion, field collection, transformation, database management, map publishing, and reporting. While each component may function effectively in isolation, the resulting workflow is frequently fragmented, difficult to reproduce, and challenging to automate.

AcuGIS Cloud addresses this challenge through a modular architecture built around deterministic workflows and operational integration. Rather than treating ingestion, synchronization, publishing, and reporting as independent activities, the platform provides a coordinated framework that allows spatial data to move through a consistent and repeatable lifecycle.

The platform combines seven primary components:

- **GeoETL** for data ingestion, inspection, transformation, and pipeline execution
- **GeoSync** for operational synchronization of field-collected data into PostGIS
- **GeoVault** for dataset lifecycle management, cataloging, versioning, and editing
- **PostGIS Service** for database provisioning and administration
- **QGIS Publisher** for publication of QGIS projects and services
- **TileServer and GeoServer integrations** for service delivery and interoperability
- **Jasper Reports and Dashboard Services** for operational reporting and analytics

Together these components provide a unified environment for managing spatial operations across the enterprise.

2.2 Operational Lifecycle

The AcuGIS Cloud architecture is organized around a six-stage operational lifecycle:

Acquire → Sync → Transform → Store → Publish → Report

Each stage represents a distinct operational concern while remaining integrated with the broader platform.

Acquire

Data acquisition is performed through GeoETL and GeoVault using a variety of supported sources, including:

- Local file uploads
- Remote URLs
- Cloud object storage
- ESRI Feature Services
- Existing PostGIS databases
- Other geospatial services and APIs

Incoming datasets are inspected automatically to determine format, schema, geometry type, coordinate reference system, and other metadata required for subsequent processing.

Sync

Operational field data is synchronized through GeoSync.

GeoSync supports synchronization workflows for:

- QFieldCloud
- Mergin Maps
- Survey123 and ArcGIS-based collection systems

Mobile and field-collected datasets are synchronized into enterprise PostGIS environments using change-based synchronization mechanisms that prioritize consistency and repeatability.

Transform

Transformation workflows are managed through GeoETL and GeoVault.

Common operations include:

- Field mapping and renaming
- Data filtering
- Schema normalization
- Geometry repair

- Coordinate reprojection
- Attribute calculations
- Spatial editing operations

Transformations are represented as explicit operations rather than undocumented manual procedures, enabling workflows to be reviewed, repeated, and automated.

Store

Persistent storage is managed through GeoVault and PostGIS Service.

GeoVault provides:

- Dataset cataloging
- Version management
- Metadata management
- Access control integration
- Dataset organization

PostGIS Service provides:

- Database provisioning
- Connection management
- Administrative operations
- Spatial database lifecycle management

Together these components establish a centralized system of record for operational spatial data.

Publish

Published outputs are delivered through multiple channels.

Depending on dataset type and operational requirements, data may be published through:

- QGIS Publisher
- TileServer
- GeoServer
- WMS services
- WFS services
- Vector tile services

- Web map applications

This architecture allows organizations to support both internal operational users and external data consumers using the same underlying datasets.

Report

Operational reporting and analytics are delivered through Jasper Reports and Dashboard Services.

Reports may be generated on demand or through scheduled execution workflows, enabling organizations to transform operational spatial data into actionable outputs for decision makers, analysts, and stakeholders.

2.3 Core Platform Components

GeoETL

GeoETL serves as the ingestion and transformation engine of the platform.

It supports data inspection, transformation design, preview workflows, deterministic execution, scheduled processing, and optional publication workflows.

A defining characteristic of GeoETL is its support for deterministic pipeline execution. Transformation logic is represented explicitly and can be reviewed, validated, repeated, and scheduled without reliance on undocumented manual procedures.

GeoSync

GeoSync provides operational synchronization between field collection systems and enterprise spatial databases.

Rather than requiring manual exports and imports, GeoSync enables datasets collected through supported mobile platforms to be synchronized into PostGIS environments through repeatable workflows.

GeoSync is designed to operationalize field data by integrating mobile collection activities directly into enterprise spatial operations.

GeoVault

GeoVault functions as the operational data catalog of AcuGIS Cloud.

It provides dataset registration, version management, metadata management, editing workflows, publication routing, and organizational controls.

GeoVault acts as the central management layer for datasets throughout their operational lifecycle.

PostGIS Service

PostGIS Service provides administrative capabilities for managing spatial databases.

Capabilities include database provisioning, connection management, schema administration, access control integration, and operational maintenance tasks.

This service separates database administration from publishing and catalog management concerns while maintaining integration with the broader platform.

Publishing Services

Publishing capabilities are provided through QGIS Publisher, TileServer, and GeoServer integrations.

These services expose operational datasets through standard geospatial protocols and web mapping technologies, enabling interoperability with desktop GIS, web applications, and third-party systems.

Reporting Services

Reporting services provide scheduled and on-demand access to operational information derived from spatial datasets.

By integrating reporting directly into the platform architecture, organizations can establish repeatable workflows that extend from data acquisition through final reporting outputs.

2.4 Architectural Principles

The AcuGIS Cloud architecture is guided by several core principles.

Deterministic Execution

Operational workflows should be reproducible.

Transformations, synchronization processes, and publication workflows should produce consistent outputs when executed repeatedly against equivalent inputs.

Human-Guided Automation

Automation should reduce operational effort without obscuring process visibility.

Users remain able to inspect, review, validate, and control workflow execution while benefiting from automation where appropriate.

Self-Hosted Deployment

Organizations retain control over infrastructure, data, and operational processes.

The platform is designed to support self-hosted deployment models while maintaining a unified operational experience.

Open Standards and Interoperability

The platform is built around widely adopted geospatial technologies and standards, including PostGIS, QGIS, GeoServer, GeoPackage, WMS, WFS, and vector tile services.

This approach minimizes vendor lock-in and simplifies integration with existing GIS investments.

Separation of Concerns

Each platform component focuses on a specific operational responsibility while integrating through clearly defined interfaces.

This modular architecture improves maintainability, scalability, and operational flexibility.

Operational Continuity

Spatial operations do not end at data ingestion.

The platform is designed to support the complete lifecycle of operational GIS activities, including acquisition, synchronization, transformation, storage, publication, and reporting.

3. GeoETL: Deterministic Spatial Pipelines

3.1 Introduction

Data acquisition and preparation remain among the most time-consuming aspects of operational GIS.

Organizations routinely obtain spatial data from a wide variety of sources, including downloaded files, cloud storage systems, APIs, enterprise databases, and third-party geospatial services. Before such data can be used operationally, it often requires inspection, cleaning, transformation, schema normalization, coordinate system conversion, and validation.

Historically, these activities have been performed through a combination of desktop GIS operations, custom scripts, spreadsheets, and ad hoc procedures. While effective for individual projects, these approaches are frequently difficult to reproduce, automate, or audit.

GeoETL was developed to address this challenge through a deterministic pipeline model for spatial data preparation.

Rather than relying on undocumented manual workflows, GeoETL represents transformation logic as explicit operations that can be inspected, reviewed, validated, executed, scheduled, and reused.

The result is a repeatable approach to spatial data preparation that supports both interactive workflows and operational automation.

3.2 Design Principles

GeoETL is built around four fundamental principles.

Explicit Transformations

Every modification applied to a dataset should be represented as an explicit operation.

Rather than embedding business logic in undocumented desktop workflows or custom scripts, GeoETL records transformation intent in a structured form that can be reviewed and repeated.

Preview Before Execution

Transformations should be visible before they are executed.

Users are able to review planned modifications, inspect schema changes, validate calculated fields, and examine sample outputs prior to import or publication.

Deterministic Execution

Equivalent inputs and equivalent operations should produce equivalent outputs.

GeoETL emphasizes reproducibility by separating transformation definition from execution and by preserving pipeline definitions for future reuse.

Human-Guided Automation

Automation should assist workflow creation without obscuring execution.

GeoETL may use natural-language assistance to simplify transformation design, but execution remains transparent, reviewable, and under user control.

3.3 Supported Data Sources

GeoETL supports ingestion from a variety of operational sources.

Typical inputs include:

- GeoPackage
- Shapefile

- GeoJSON
- CSV
- Microsoft Excel
- Cloud Optimized GeoTIFF (COG)
- PostGIS
- Remote URLs
- Cloud object storage
- ESRI Feature Services
- REST APIs

Upon acquisition, datasets are inspected automatically to identify:

- Geometry type
- Coordinate reference system
- Schema structure
- Field definitions
- Attribute statistics
- Sample records

This inspection process provides the foundation for subsequent transformation and validation workflows.

3.4 Transformation Design

GeoETL supports both structured and natural-language approaches to transformation design.

Users may define transformations directly through the interface or describe intended changes using natural-language instructions.

Examples include:

- Rename parcel_id to parcel_number
- Remove temporary fields
- Filter inactive records
- Calculate new attributes
- Normalize schema definitions
- Reproject coordinates

Natural-language requests are translated into explicit transformation operations that remain fully visible and editable before execution.

This approach combines the accessibility of natural-language interaction with the predictability of deterministic processing.

Figure 2. GeoETL transformation design interface

GeoETL converts user-defined transformation requests into explicit operations that can be reviewed, modified, and validated before execution.

3.5 Deterministic Pipeline Model

The defining characteristic of GeoETL is its deterministic pipeline architecture.

A pipeline represents a complete transformation workflow, including:

- Source definition
- Inspection metadata
- Transformation operations
- Validation rules
- Destination configuration
- Publication options
- Scheduling metadata

Pipelines are stored as structured definitions rather than procedural scripts.

This enables workflows to be:

- Repeated
- Shared
- Scheduled
- Audited
- Versioned
- Automated

A pipeline created today can be executed against future versions of the same dataset while preserving transformation consistency.

This capability is particularly valuable for recurring operational imports where source data is updated regularly but processing requirements remain unchanged.

3.6 Transformation Operations

GeoETL supports a broad range of transformation operations commonly required in spatial data workflows.

Examples include:

Schema Operations

- Add field
- Remove field
- Rename field
- Reorder fields
- Change data types

Attribute Operations

- Compute derived values
- String manipulation
- Numeric calculations
- Conditional expressions
- Value normalization

Spatial Operations

- Coordinate reprojection
- Geometry validation
- Geometry repair
- CRS normalization
- Spatial filtering

Dataset Operations

- Row filtering
- Record selection
- Data cleansing
- Null handling
- Duplicate management

Operations may be combined into complex workflows while remaining individually visible and editable.

3.7 Preview and Validation

Before execution, GeoETL generates a transformation preview that summarizes the expected effects of the pipeline.

Preview information may include:

- Planned operations
- Schema modifications
- Added fields
- Removed fields
- Sample value changes
- Validation results

This review step serves as a quality-control checkpoint between transformation design and execution.

By exposing changes before import, GeoETL reduces the risk of unintended modifications while improving transparency and auditability.

Figure 3. GeoETL transformation preview

Transformation previews summarize schema changes, calculated fields, and representative data modifications before execution.

3.8 Import and Execution

Once a transformation has been reviewed and approved, GeoETL executes the pipeline against the selected source dataset.

Supported destinations may include:

- PostGIS
- GeoPackage
- GeoJSON
- Other supported spatial formats

During execution, GeoETL performs any required validation, schema normalization, geometry repair, or coordinate transformations before loading the dataset into its destination.

Users may execute a workflow immediately or save the pipeline for future reuse.

Figure 4. GeoETL import and execution workflow

Approved transformations can be executed directly or saved as reusable operational pipelines.

3.9 Scheduling and Operational Automation

Many GIS workflows involve recurring imports from sources that are updated regularly.

Examples include:

- Municipal open-data portals
- Environmental monitoring feeds
- Utility asset inventories
- Census and demographic updates
- Field-collected datasets

GeoETL supports scheduled execution of saved pipelines, allowing organizations to automate repetitive ingestion and transformation workflows while maintaining deterministic processing behavior.

Because scheduling operates on saved pipeline definitions rather than ad hoc procedures, organizations can automate workflows without sacrificing reproducibility or transparency.

3.10 Integration with the AcuGIS Cloud Platform

GeoETL serves as the primary acquisition and transformation engine within AcuGIS Cloud.

After execution, datasets may flow into:

- GeoVault for cataloging and lifecycle management

- PostGIS Service for database storage
- GeoServer for publication
- QGIS Publisher for map deployment
- GeoSync workflows for operational integration
- Jasper Reports and dashboards for downstream analysis

This integration enables a continuous operational lifecycle spanning acquisition, synchronization, transformation, storage, publication, and reporting.

GeoETL therefore functions not merely as an import utility, but as the foundation of deterministic spatial operations across the AcuGIS Cloud platform.

3.11 Summary

GeoETL addresses one of the most persistent challenges in GIS: transforming spatial data in a repeatable and auditable manner.

By combining inspection, preview, deterministic execution, reusable pipelines, and operational automation, GeoETL enables organizations to move beyond manual data preparation workflows and toward reproducible spatial operations.

Within AcuGIS Cloud, GeoETL provides the foundation for consistent acquisition and transformation workflows that integrate directly with synchronization, storage, publishing, and reporting systems.

4. GeoSync: Operational Field Synchronization

4.1 Introduction

Field data collection has become a critical component of modern GIS operations.

Organizations increasingly rely on mobile applications to collect inspections, environmental observations, utility asset updates, infrastructure inventories, and operational field data. While mobile collection platforms have become increasingly capable, synchronization of collected information into enterprise spatial databases remains a persistent challenge.

Many organizations continue to rely on manual exports, periodic imports, custom synchronization scripts, or project-specific workflows that are difficult to standardize and maintain.

The problem is not data collection itself.

The problem is operationalizing field data.

GeoSync was developed to address this challenge by providing a repeatable and deterministic synchronization layer between mobile GIS platforms and enterprise PostGIS environments.

Within AcuGIS Cloud, GeoSync represents the **Sync** stage of the operational lifecycle:

Acquire → Sync → Transform → Store → Publish → Report

Its purpose is to move field-collected spatial data into operational systems reliably, predictably, and with minimal manual intervention.

4.2 The Mobile GIS Challenge

Mobile GIS platforms are optimized for data collection in disconnected and operational environments.

However, field workflows introduce several challenges:

- Offline editing
- Intermittent connectivity
- Multiple collection platforms
- Dataset duplication
- Import/export complexity
- Inconsistent update workflows
- Difficult integration with enterprise databases

Organizations frequently encounter situations where field crews successfully collect data, yet operational teams struggle to integrate those updates into production databases.

The result is a disconnect between collection and operations.

Data may exist on mobile devices or synchronization platforms but remain unavailable for publication, reporting, analysis, or operational decision-making.

GeoSync addresses this problem by establishing a consistent synchronization workflow between supported field systems and enterprise PostGIS databases.

4.3 Supported Synchronization Platforms

GeoSync supports synchronization workflows for several widely used field collection environments.

QFieldCloud

GeoSync can synchronize projects managed through QFieldCloud into operational PostGIS environments.

This enables organizations to maintain QGIS-based field collection workflows while integrating collected data directly into enterprise spatial systems.

Mergin Maps

GeoSync supports synchronization of Mergin Maps projects using the Mergin synchronization ecosystem.

Changes collected through Mergin-based workflows can be incorporated into PostGIS through repeatable synchronization processes.

Survey123

GeoSync supports Survey123 and ArcGIS-based collection workflows, allowing organizations to integrate ESRI field collection activities with PostGIS-based operational infrastructure.

Self-Hosted Workflows

Organizations may also configure self-hosted synchronization providers and existing service endpoints where operational requirements demand additional flexibility.

Figure 5. GeoSync provider configuration

GeoSync supports synchronization workflows across multiple field collection platforms through a unified configuration interface.

4.4 Synchronization Architecture

GeoSync is intentionally designed as an operational synchronization service rather than a generalized replication platform.

Its objective is straightforward:

1. Connect to a supported provider
2. Detect project changes

3. Retrieve updated data
4. Apply changes to PostGIS
5. Maintain operational consistency

This model reduces complexity while providing predictable synchronization behavior.

At a high level, synchronization follows the pattern:

```
Field Platform
       ↓
GeoPackage / Project Data
       ↓
GeoSync
       ↓
Synchronization Engine
       ↓
PostGIS
```

Depending on the source platform, GeoSync may utilize provider-specific synchronization engines and workflows while presenting a consistent operational interface to administrators.

This abstraction allows organizations to work with multiple field collection technologies without requiring separate operational procedures for each system.

4.5 Project-Based Synchronization

Synchronization is organized around projects.

Each project defines:

- Source provider
- Project identifier
- Synchronization interval
- Target database
- Target schema
- Operational settings

Administrators may manage multiple synchronization projects simultaneously through a centralized interface.

This approach allows organizations to support multiple teams, departments, or field initiatives while maintaining consistent synchronization practices.

Figure 6. GeoSync project configuration

Synchronization schedules and project settings are managed centrally through project-level configuration.

4.6 Automated Synchronization

A core objective of GeoSync is reducing manual operational effort.

Synchronization projects can be configured to execute automatically at defined intervals.

Typical synchronization frequencies include:

- Every few minutes
- Hourly
- Daily
- Project-specific schedules

During execution, GeoSync monitors configured projects, retrieves updates, and applies changes to the target PostGIS environment.

This process enables organizations to maintain current operational datasets without requiring repeated manual imports.

By automating synchronization activities, field-collected information becomes available to downstream systems more quickly and more consistently.

4.7 Change-Based Synchronization

GeoSync emphasizes change-based synchronization rather than wholesale dataset replacement.

This approach offers several advantages:

- Reduced processing overhead

- Faster synchronization cycles
- Lower bandwidth requirements
- Improved operational efficiency
- Better support for frequently updated datasets

Only the changes required to bring the target system into alignment are applied.

This model is particularly valuable for operational environments where datasets are updated continuously and synchronization occurs frequently throughout the day.

4.8 Integration with PostGIS

PostGIS serves as the operational destination for synchronized data.

Once synchronized, datasets become available to the broader AcuGIS Cloud ecosystem.

This includes:

- GeoVault catalog management
- GeoETL transformation workflows
- QGIS Publisher deployment
- GeoServer publication
- Reporting and dashboards
- Custom applications and APIs

By synchronizing directly into PostGIS, GeoSync avoids creating isolated data silos and instead contributes to a unified operational data environment.

4.9 Operational Workflows

GeoSync is intended to support a variety of operational scenarios.

Field Inspections

Inspection teams collect observations using mobile applications.

GeoSync transfers updates into PostGIS, where information becomes available for mapping, reporting, and analysis.

Environmental Monitoring

Field personnel collect environmental measurements and observations.

Updates are synchronized automatically and made available to dashboards and reporting systems.

Asset Management

Utility and infrastructure organizations can synchronize field updates into enterprise spatial databases, reducing delays between collection and operational visibility.

Community and Survey Programs

Organizations conducting recurring surveys may synchronize collected information into centralized systems for publication and reporting.

In each case, GeoSync acts as the bridge between collection activities and enterprise operations.

4.10 Integration with the AcuGIS Cloud Lifecycle

GeoSync occupies a distinct role within the AcuGIS Cloud architecture.

While GeoETL focuses on acquisition and transformation, GeoSync focuses on operational synchronization.

Together the two systems address different aspects of the spatial data lifecycle:

GeoETL

- Imports external datasets
- Performs transformation workflows
- Executes deterministic pipelines
- Supports recurring operational imports

GeoSync

- Synchronizes field-collected data
- Maintains operational datasets
- Automates project updates
- Connects mobile workflows to PostGIS

These capabilities complement one another and together establish a complete operational pathway from data acquisition through publication and reporting.

4.11 Architectural Tradeoffs

GeoSync intentionally prioritizes operational simplicity and predictability.

The platform is designed to provide:

- Deterministic synchronization behavior
- Repeatable operational workflows
- Centralized project management
- Direct integration with PostGIS
- Support for multiple collection platforms

Rather than functioning as a generalized data replication framework, GeoSync focuses on the specific challenge of moving field-collected spatial information into enterprise GIS environments.

This design decision reduces operational complexity while improving reliability and maintainability.

4.12 Summary

GeoSync addresses one of the most common challenges in operational GIS: integrating field-collected data into enterprise spatial systems.

By supporting multiple collection platforms, automating synchronization workflows, and integrating directly with PostGIS, GeoSync enables organizations to operationalize mobile GIS activities through consistent and repeatable processes.

Within AcuGIS Cloud, GeoSync provides the synchronization layer that connects field collection workflows to transformation, storage, publication, and reporting systems.

Together with GeoETL, it forms the foundation of a unified spatial operations platform built around deterministic workflows and operational continuity.

5. GeoVault: Dataset Lifecycle Management

5.1 Introduction

Spatial data management extends beyond storage.

Organizations must not only store datasets, but also organize, version, edit, publish, discover, and maintain them throughout their operational lifecycle. As the number of datasets grows, managing these activities through shared folders, isolated databases, or ad hoc publishing workflows becomes increasingly difficult.

GeoVault was developed to address this challenge by providing a centralized platform for dataset lifecycle management.

Within AcuGIS Cloud, GeoVault serves as the primary operational catalog for spatial data and occupies the **Store** stage of the platform lifecycle:

Acquire → Sync → Transform → Store → Publish → Report

GeoVault provides a unified environment where datasets can be organized, versioned, edited, transformed, published, and integrated with downstream services.

Rather than functioning solely as a repository, GeoVault acts as the operational hub through which datasets move from acquisition to publication.

5.2 Dataset-Centric Architecture

Traditional GIS systems often organize workflows around servers, databases, or applications.

GeoVault adopts a dataset-centric model.

Each dataset becomes a managed operational asset with associated:

- Metadata
- Versions
- Publication settings
- Transformation history
- Editing capabilities
- Access controls
- Service integrations

This approach provides users with a consistent experience regardless of how data was originally acquired.

Whether a dataset originated from:

- GeoETL imports
- GeoSync synchronization
- PostGIS databases
- Uploaded files
- External services

it becomes a managed asset within the GeoVault catalog.

5.3 Organizational Model

GeoVault organizes datasets using a folder-based structure designed to support operational workflows.

Folders may represent:

- Projects
- Departments
- Geographic regions
- Programs
- Clients
- Operational teams

This model provides a familiar organizational framework while remaining independent of underlying storage technologies.

Datasets can be grouped according to business requirements rather than technical constraints.

Figure 7. GeoVault operational catalog

GeoVault organizes datasets through a folder-based catalog while exposing publishing and database services through a unified management interface.

This organizational model allows both technical and non-technical users to locate and manage spatial assets without requiring direct access to underlying databases or file systems.

5.4 Dataset Lifecycle Management

A core objective of GeoVault is managing datasets throughout their operational lifecycle.

Typical lifecycle activities include:

Registration

Datasets enter GeoVault through:

- GeoETL imports

- GeoSync synchronization
- File uploads
- PostGIS imports
- Administrative workflows

Maintenance

Once registered, datasets may be:

- Edited
- Updated
- Republished
- Versioned
- Archived
- Restored

Publication

Datasets can be routed into publication workflows through integrated publishing services.

Retirement

Datasets that are no longer operationally active may be archived while preserving historical records and associated metadata.

This lifecycle-oriented approach provides operational continuity while reducing the need for separate management systems.

5.5 Version Management and Lineage

One of GeoVault's most significant capabilities is dataset version management.

Many GIS workflows rely on direct replacement of files or database tables, making it difficult to understand how a dataset evolved over time.

GeoVault addresses this challenge through managed dataset versions.

Each version may preserve:

- Creation date
- Source information
- Transformation history
- Publication status
- Operational lineage

Users can review historical versions, compare changes, and restore previous states when necessary.

Figure 8. GeoVault version management

Dataset versions preserve lineage, transformation history, publication status, and rollback capabilities.

This capability is particularly valuable in operational environments where datasets are updated regularly and historical traceability is required.

5.6 Lineage and Traceability

Operational GIS workflows frequently involve multiple processing steps.

A dataset may originate from an external source, pass through transformation workflows, undergo editing, and eventually be published through multiple services.

GeoVault maintains visibility into this lifecycle through lineage tracking.

Lineage information helps answer questions such as:

- Where did this dataset originate?
- Which transformations were applied?
- Which version is currently published?
- When was the dataset last updated?
- Which datasets depend on this version?

This improves transparency while simplifying operational management.

Lineage tracking also supports auditing and quality assurance activities by providing visibility into how datasets evolve over time.

5.7 Browser-Based Spatial Editing

GeoVault provides browser-based editing capabilities for operational datasets.

Users can interact directly with spatial features through integrated mapping interfaces and attribute forms.

Supported activities include:

- Feature editing
- Attribute updates
- Record creation
- Record deletion
- Geometry modification
- Operational maintenance workflows

Figure 9. GeoVault spatial editing

Browser-based editing combines map interaction with attribute management to support operational maintenance workflows.

This capability enables organizations to perform many common maintenance activities without requiring desktop GIS software.

For operational teams, browser-based editing reduces deployment complexity while improving accessibility.

5.8 Integration with GeoETL

GeoVault and GeoETL are closely integrated.

GeoETL focuses on acquisition and transformation, while GeoVault focuses on lifecycle management.

After import through GeoETL:

- Datasets are cataloged within GeoVault
- Metadata becomes available
- Versions are tracked
- Publication workflows become available
- Editing workflows may be enabled

This integration establishes a continuous pathway from ingestion to operational management.

Rather than ending at import, GeoETL workflows transition naturally into ongoing dataset administration.

5.9 Publication Routing

GeoVault serves as a bridge between managed datasets and publication services.

Users can initiate publication workflows directly from dataset management interfaces.

Depending on operational requirements, datasets may be published through:

- QGIS Publisher
- GeoServer
- Tile services
- Web applications
- Reporting workflows

This integration reduces friction between data management and service delivery.

Datasets do not need to be exported or manually transferred between systems before publication.

5.10 Operational Metadata

Metadata plays a critical role in dataset discovery and management.

GeoVault maintains operational metadata associated with managed datasets.

Typical metadata includes:

- Dataset name
- Description
- Creation date
- Version information
- Publication status
- Ownership information

- Organizational context

This information assists users in locating datasets, understanding operational status, and evaluating suitability for specific purposes.

Metadata also provides the foundation for catalog and discovery functionality across the platform.

5.11 Archive and Retention Management

Operational environments often require datasets to be retained even after active use has ended.

GeoVault supports archive workflows that allow datasets to be removed from active operations while preserving their history.

Archived datasets remain available for:

- Historical reference
- Audit requirements
- Restoration
- Reuse

This capability helps organizations manage growth while maintaining access to historical information.

5.12 Integration with the AcuGIS Cloud Platform

GeoVault functions as the operational center of the AcuGIS Cloud ecosystem.

It connects directly with:

GeoETL

For ingestion and transformation workflows.

GeoSync

For synchronized operational datasets.

PostGIS Service

For database storage and administration.

QGIS Publisher

For map publication and deployment.

GeoServer

For standards-based service delivery.

Jasper Reports and Dashboards

For operational reporting and analytics.

Because of these integrations, GeoVault acts as the primary management layer through which datasets move across the platform.

5.13 Architectural Principles

GeoVault is designed around several key principles.

Dataset-Centric Management

Datasets should be treated as managed operational assets rather than isolated files or database tables.

Operational Continuity

Management workflows should extend beyond ingestion and support the complete operational lifecycle.

Traceability

Versioning and lineage information should remain available throughout the life of a dataset.

Accessibility

Common maintenance activities should be available through browser-based interfaces.

Integration

Dataset management should connect naturally with transformation, publication, synchronization, and reporting workflows.

These principles allow GeoVault to function as more than a repository while remaining approachable to both technical and operational users.

5.14 Summary

GeoVault provides the dataset lifecycle management layer of AcuGIS Cloud.

Through cataloging, versioning, lineage tracking, browser-based editing, publication routing, and operational metadata management, GeoVault enables organizations to manage spatial data as operational assets rather than isolated files or database tables.

Together with GeoETL and GeoSync, GeoVault establishes a continuous workflow that spans acquisition, synchronization, transformation, storage, publication, and reporting.

As the central operational catalog of AcuGIS Cloud, GeoVault provides the foundation upon which publishing, reporting, and broader spatial operations are built.

6. PostGIS Service: Spatial Database Lifecycle Management

6.1 Introduction

Spatial databases are the operational foundation of most enterprise GIS environments.

While map servers, dashboards, reporting systems, and data catalogs often receive significant attention, the underlying spatial database remains the system of record upon which these services depend.

As GIS deployments grow, organizations frequently encounter challenges related to database provisioning, connection management, access control, schema administration, backup procedures, and operational governance.

Historically, these activities have been performed through direct database administration tools or custom operational procedures that are disconnected from broader GIS workflows.

PostGIS Service was developed to address this challenge by providing a dedicated database lifecycle management layer within AcuGIS Cloud.

Rather than treating spatial databases as isolated infrastructure components, PostGIS Service integrates database administration directly into the broader spatial operations platform.

Within the AcuGIS Cloud lifecycle, PostGIS Service supports the **Store** stage by managing the databases that serve as the operational foundation for acquisition, synchronization, publishing, and reporting workflows.

6.2 The Database Management Challenge

Many GIS organizations maintain multiple databases across development, testing, staging, and production environments.

As deployments expand, administrative complexity increases.

Common challenges include:

- Provisioning new databases
- Managing connection information
- Creating and maintaining schemas
- Controlling user access
- Supporting operational teams
- Coordinating publishing systems
- Managing backup and recovery workflows

These responsibilities often fall to a small number of database administrators or GIS specialists.

As a result, routine operational tasks can become bottlenecks that slow project delivery and increase administrative overhead.

AcuGIS Cloud addresses this challenge by providing a dedicated service focused specifically on the management of PostGIS environments.

6.3 Separation of Responsibilities

A key architectural principle of AcuGIS Cloud is separation of concerns.

Each platform component focuses on a distinct operational responsibility.

Within this model:

GeoETL

Manages acquisition and transformation workflows.

GeoSync

Manages field synchronization workflows.

GeoVault

Manages datasets and lifecycle operations.

PostGIS Service

Manages databases.

This distinction is important.

GeoVault treats datasets as managed operational assets.

PostGIS Service treats databases as managed infrastructure resources.

Although these systems are closely integrated, they address different operational requirements.

This separation improves maintainability while reducing complexity across the platform.

6.4 Database Provisioning

Creating and configuring new spatial databases is a common requirement in GIS operations.

PostGIS Service provides administrative workflows that simplify provisioning activities.

Typical provisioning tasks include:

- Creating databases
- Enabling PostGIS extensions
- Managing schemas
- Configuring access
- Registering database environments

These capabilities allow organizations to establish new operational environments without requiring manual database configuration for every project.

By integrating provisioning workflows into the platform, AcuGIS Cloud reduces administrative effort while promoting consistency across environments.

6.5 Connection Management

Modern GIS environments often interact with multiple databases simultaneously.

Examples include:

- Development databases
- Production databases
- External client systems
- Department-specific environments
- Data warehouse integrations

Managing these connections manually can become increasingly difficult as infrastructure grows.

PostGIS Service centralizes connection management and provides a consistent mechanism for registering and administering database resources.

This capability improves visibility while reducing duplication of configuration information across the platform.

Connection definitions can then be reused by:

- GeoETL
- GeoVault

- QGIS Publisher
- Reporting systems
- Administrative workflows

6.6 Schema Administration

Schemas provide an important organizational layer within enterprise PostGIS environments.

Organizations frequently separate operational datasets according to:

- Projects
- Departments
- Clients
- Programs
- Applications

PostGIS Service supports administrative workflows that simplify schema management and improve organizational consistency.

By exposing schema information through the platform, operational users gain improved visibility into database structure without requiring direct administrative access to database servers.

This approach helps bridge the gap between GIS operations and database administration.

6.7 Access Control and Governance

Database governance becomes increasingly important as GIS deployments expand.

Organizations require mechanisms for controlling:

- Database access
- User permissions
- Role assignments
- Administrative privileges
- Operational responsibilities

PostGIS Service supports governance workflows by integrating database administration into the broader authentication and authorization model of AcuGIS Cloud.

This enables organizations to manage operational access while maintaining consistency across platform services.

Rather than implementing separate security models for each component, AcuGIS Cloud provides a unified approach to platform administration.

6.8 Operational Administration

Beyond initial provisioning, databases require ongoing operational management.

Typical administrative activities include:

- Environment maintenance
- Database cloning
- Configuration updates
- User management
- Access reviews
- Operational troubleshooting

PostGIS Service provides administrative capabilities designed specifically for spatial database environments.

These workflows reduce dependence on external tools while allowing organizations to perform routine operational tasks through a consistent interface.

As a result, administrative activities become more accessible to GIS operations teams while remaining compatible with established database management practices.

6.9 Integration with GeoETL

GeoETL frequently requires access to PostGIS environments for data import and transformation workflows.

PostGIS Service provides the database management foundation that supports these activities.

Typical interactions include:

- Import destinations
- Database selection
- Schema targeting
- Connection reuse
- Operational storage

This integration allows GeoETL to focus on transformation and execution while PostGIS Service manages the underlying database environment.

The result is a cleaner architectural separation between data processing and database administration.

6.10 Integration with GeoSync

GeoSync synchronizes field-collected information directly into PostGIS environments.

PostGIS Service provides the operational infrastructure required to support these synchronization workflows.

Through this integration:

- Synchronization targets can be managed centrally
- Database environments remain visible across the platform
- Operational governance remains consistent
- Database resources can be reused across projects

This relationship enables field-collected information to flow into managed operational environments rather than isolated project databases.

6.11 Integration with Publishing Services

Publishing systems depend on reliable access to spatial databases.

QGIS Publisher, GeoServer, reporting services, dashboards, and custom applications all consume information stored within PostGIS environments.

PostGIS Service supports these workflows by providing:

- Database administration
- Connection visibility
- Operational consistency
- Infrastructure management

This architecture allows publishing systems to focus on service delivery while database administration remains centralized.

The result is improved maintainability and reduced duplication of operational effort.

6.12 Architectural Benefits

The inclusion of a dedicated database lifecycle service provides several advantages.

Administrative Consistency

Database management activities are standardized across projects and environments.

Reduced Operational Complexity

Routine database administration becomes more accessible to operational teams.

Improved Governance

Database resources remain visible and manageable within a centralized platform.

Better Integration

Database administration becomes part of the broader spatial operations lifecycle rather than a disconnected infrastructure activity.

Scalability

Organizations can manage growing numbers of databases and environments through a consistent operational model.

Together these benefits support the long-term maintainability of enterprise GIS deployments.

6.13 Role Within AcuGIS Cloud

PostGIS Service serves as the infrastructure foundation of AcuGIS Cloud.

While users may interact primarily with GeoETL, GeoSync, GeoVault, publishing services, or reporting tools, many of these activities ultimately depend on managed PostGIS environments.

PostGIS Service provides the operational layer that connects platform workflows to the databases that store and serve spatial information.

It enables AcuGIS Cloud to treat database administration as an integrated component of spatial operations rather than an external responsibility.

6.14 Summary

PostGIS Service provides the database lifecycle management layer of AcuGIS Cloud.

Through provisioning, connection management, schema administration, governance, and operational administration, it enables organizations to manage spatial database environments as part of a unified spatial operations platform.

By separating database management from dataset management while maintaining close integration with GeoETL, GeoSync, GeoVault, publishing services, and reporting systems, PostGIS Service establishes a scalable foundation for enterprise GIS operations.

Within AcuGIS Cloud, it serves as the infrastructure layer that supports the storage, management, and delivery of operational spatial data.

7. Publishing Architecture

7.1 Introduction

The ultimate objective of most GIS workflows is not data collection or storage.

It is delivery.

Organizations invest in data acquisition, synchronization, transformation, and management so that information can be shared with decision makers, operational staff, analysts, stakeholders, and the public.

Publishing therefore represents the stage at which spatial information becomes operationally useful.

Within AcuGIS Cloud, publishing is not treated as a separate activity performed after data management. Instead, publication capabilities are integrated directly into the broader spatial operations lifecycle.

This architecture enables datasets to move seamlessly from acquisition and synchronization workflows into maps, services, applications, dashboards, and reports.

Within the platform lifecycle, publishing represents the **Publish** stage:

Acquire → Sync → Transform → Store → Publish → Report

7.2 Publishing Philosophy

Many GIS platforms require users to navigate multiple disconnected systems before data becomes available to end users.

Typical workflows often involve:

- Loading data into a database
- Configuring map services
- Creating web applications
- Managing cache generation
- Building reports
- Maintaining service endpoints

These activities frequently require different tools, administrative procedures, and technical expertise.

AcuGIS Cloud adopts a different approach.

Publishing is integrated directly into operational workflows.

Datasets managed within GeoVault can move directly into publishing environments without requiring manual transfers between systems.

This integration reduces operational complexity while accelerating delivery.

7.3 Publishing Components

The publishing architecture combines several specialized services that work together to deliver spatial information.

Core components include:

QGIS Publisher

Project-based publishing and application deployment.

GeoServer

Standards-based OGC service delivery.

Tile Services

Vector and raster tile publication.

Web Map Applications

Interactive browser-based GIS applications.

Reporting Services

Operational reporting integrated directly into published applications.

Together these components provide multiple pathways for delivering spatial information while remaining connected through a unified management experience.

7.4 QGIS Publisher

QGIS Publisher serves as the primary publishing environment within AcuGIS Cloud.

It enables organizations to deploy QGIS projects as web-accessible applications while preserving project structure, styling, symbology, and operational workflows.

QGIS Publisher supports:

- QGIS project deployment
- Project version management

- Symbology administration
- Cache management
- Reporting integration
- Operational updates

Published projects can be managed through a centralized interface that provides visibility into deployment status and associated services.

Figure 10. QGIS Publisher project management

QGIS Publisher provides centralized management of published projects, versions, cache operations, symbology, and reporting integrations.

This approach allows organizations to leverage existing QGIS workflows while extending them into operational web environments.

7.5 Raster and Vector Publishing

Modern GIS environments must support a variety of spatial data types.

AcuGIS Cloud supports publication of both raster and vector datasets.

Raster Publishing

Typical raster workflows include:

- GeoTIFF
- Cloud Optimized GeoTIFF (COG)
- Orthophotography
- Environmental surfaces
- Elevation models

Raster datasets may be published through QGIS Publisher, GeoServer, or tile-based services depending on operational requirements.

Vector Publishing

Vector workflows support:

- Points
- Lines

- Polygons
- Operational datasets
- Administrative boundaries
- Asset inventories

Vector datasets can be delivered through traditional map services, vector tile services, or application-specific APIs.

This flexibility allows organizations to select delivery mechanisms appropriate for their users and performance requirements.

7.6 Service-Oriented Publishing

Many organizations require standards-based service delivery for interoperability with desktop GIS, external applications, and enterprise systems.

AcuGIS Cloud supports service-oriented publishing through integrated service infrastructure.

Typical service types include:

- WMS
- WFS
- WMTS
- Vector tiles
- REST endpoints

These services enable published datasets to be consumed by:

- QGIS
- ArcGIS
- Web applications
- Mobile applications
- Reporting systems
- Custom software

By supporting both application-based and service-based publishing models, AcuGIS Cloud accommodates a broad range of deployment scenarios.

7.7 Integrated Cache Management

Performance becomes increasingly important as datasets and user populations grow.

To support scalable deployments, publishing workflows include integrated cache management capabilities.

Cache operations may include:

- Tile generation
- Tile seeding
- Cache updates
- Cache refresh workflows

These capabilities improve responsiveness while reducing server workload for frequently accessed datasets.

Cache management is integrated directly into publishing workflows, allowing administrators to manage performance without requiring separate operational procedures.

7.8 Project Versioning

Published applications often evolve over time.

Organizations may update:

- Data sources
- Symbology
- Layouts
- Reports
- Operational workflows

QGIS Publisher supports project version management to help organizations maintain control over these changes.

Versioning capabilities provide:

- Historical project snapshots
- Rollback workflows
- Controlled updates

- Operational continuity

This capability complements the dataset versioning functionality provided by GeoVault and extends lifecycle management into the publication layer.

Together these systems support traceability from source data through published applications.

7.9 Reporting Integration

One of the distinguishing characteristics of AcuGIS Cloud is the integration of reporting directly into published applications.

Traditional GIS workflows frequently separate mapping and reporting into different systems.

Users often navigate between:

- GIS applications
- Reporting portals
- Dashboard environments

to obtain a complete operational picture.

AcuGIS Cloud integrates these activities.

Published applications can expose reporting workflows directly within the user experience.

Figure 11. Integrated mapping and reporting

Published applications can combine interactive mapping, feature information, and operational reports within a unified workflow.

This approach reduces context switching while improving operational efficiency.

Users can move directly from spatial exploration to analytical reporting without leaving the application.

7.10 GeoServer Integration

GeoServer provides standards-based publishing capabilities for organizations that require interoperable spatial services.

Within AcuGIS Cloud, GeoServer can be used to:

- Publish vector layers
- Publish raster layers
- Deliver OGC services
- Support external GIS clients
- Provide service-oriented architectures

GeoServer complements QGIS Publisher by focusing on standards-based service delivery while allowing organizations to maintain centralized dataset management through GeoVault.

This integration allows organizations to combine operational management with open standards and interoperability.

7.11 Operational Publishing Workflow

Publishing typically follows a predictable operational sequence.

Step 1

Acquire or synchronize data through GeoETL or GeoSync.

Step 2

Manage datasets through GeoVault.

Step 3

Store operational information within PostGIS.

Step 4

Publish datasets through QGIS Publisher, GeoServer, or tile services.

Step 5

Expose information through applications, reports, dashboards, and APIs.

This workflow illustrates how publishing is integrated into the broader AcuGIS Cloud lifecycle rather than functioning as an isolated activity.

7.12 Integration with GeoVault

Publishing workflows are closely integrated with GeoVault.

From within GeoVault, users can:

- Review datasets
- Manage versions
- Track lineage
- Initiate publication
- Update services

This integration eliminates many of the manual steps traditionally associated with GIS publishing.

Rather than exporting files and configuring services separately, users can move directly from dataset management to operational delivery.

7.13 Integration with Reporting

Publishing and reporting are tightly coupled within AcuGIS Cloud.

Published applications can interact with:

- Jasper Reports
- Dashboards
- Operational analytics
- Custom reporting workflows

This integration allows organizations to present both spatial and analytical information through a unified operational environment.

Rather than viewing maps and reports as separate deliverables, AcuGIS Cloud treats them as complementary components of a single operational workflow.

7.14 Architectural Benefits

The publishing architecture provides several operational advantages.

Unified Workflow

Publishing is integrated into acquisition, synchronization, and management workflows.

Multiple Delivery Models

Organizations can support applications, services, tiles, reports, and APIs through a common platform.

Operational Consistency

Dataset management, publishing, and reporting remain connected through shared lifecycle processes.

Scalability

Integrated caching and service-oriented publishing support growing workloads and user populations.

Interoperability

Support for standards-based services enables integration with external systems and applications.

Together these capabilities allow organizations to deliver spatial information efficiently while maintaining operational control.

7.15 Summary

The publishing architecture of AcuGIS Cloud provides the mechanisms through which spatial information becomes operationally accessible.

By integrating QGIS Publisher, GeoServer, tile services, web applications, and reporting workflows into a unified platform, AcuGIS Cloud enables organizations to move seamlessly from managed datasets to operational delivery.

Publishing is not treated as a separate activity performed after data management. Instead, it is a natural continuation of the broader spatial operations lifecycle.

This integration enables organizations to deliver maps, services, reports, and analytical products through a consistent and scalable operational architecture.

8. Reporting and Dashboards

8.1 Introduction

The final objective of most GIS workflows is decision support.

Organizations collect, synchronize, transform, manage, and publish spatial information so that users can better understand conditions, monitor operations, evaluate performance, and make informed decisions.

Maps play a critical role in this process, but maps alone are often insufficient.

Operational users frequently require:

- Tabular summaries
- Statistical analysis
- Key performance indicators
- Scheduled reports
- Dashboard visualizations
- Regulatory reporting
- Executive summaries

For this reason, AcuGIS Cloud treats reporting as a first-class component of the spatial operations lifecycle.

Within the platform architecture, reporting represents the final stage of the workflow:

Acquire → Sync → Transform → Store → Publish → Report

Reporting transforms operational spatial information into actionable business intelligence.

8.2 The GIS Reporting Challenge

Many organizations operate separate GIS and reporting environments.

Spatial information may reside within:

- GIS databases
- Map services
- Field collection systems

while reporting activities occur within:

- Business intelligence platforms
- Spreadsheet workflows
- Data warehouses

- Custom reporting tools

This separation introduces several challenges.

Duplicate Data Pipelines

Organizations often create secondary workflows solely for reporting purposes.

Delayed Information

Reports may rely on exported datasets rather than operational information.

Fragmented User Experience

Users must move between mapping systems and reporting systems to answer operational questions.

Increased Maintenance

Multiple reporting and GIS environments require separate administration and support.

AcuGIS Cloud addresses these challenges through direct integration of reporting and spatial operations.

8.3 Reporting Architecture

The reporting architecture is designed around a simple principle:

Reports should consume the same operational datasets used by maps, services, and applications.

Rather than creating separate reporting databases or export workflows, AcuGIS Cloud enables reports to operate directly against managed operational data sources.

This architecture provides several advantages:

- Reduced duplication
- Improved consistency
- Faster updates
- Simplified administration
- Better integration with GIS workflows

Reports therefore become a continuation of operational spatial workflows rather than a separate technology stack.

8.4 Jasper-Based Reporting

AcuGIS Cloud incorporates Jasper Reports as its primary reporting engine.

Jasper provides a mature and widely adopted platform for operational reporting and document generation.

Supported report capabilities include:

- Tabular reporting
- Summary reporting
- Statistical analysis
- Parameterized reports
- PDF generation
- Spreadsheet exports
- Scheduled execution
- Automated distribution

By integrating Jasper directly into the platform, AcuGIS Cloud enables organizations to combine spatial information with enterprise reporting workflows.

Figure 12. Parameterized operational reporting

Reports support user-defined parameters, interactive filtering, and multiple export formats.

8.5 Parameterized Reports

Operational users frequently need to answer different questions using the same report design.

Parameterized reporting allows users to customize report execution without modifying report templates.

Examples include:

- Date ranges
- Geographic areas
- Administrative boundaries
- Asset categories
- Operational thresholds
- Scenario assumptions

Parameters can be supplied through interactive forms, URLs, scheduled executions, or application workflows.

This capability enables a single report design to support a wide variety of operational requirements.

8.6 Scheduled Reporting

Many reporting activities occur on recurring schedules.

Examples include:

- Daily operational summaries
- Weekly status reports
- Monthly compliance reports
- Quarterly performance reviews
- Executive briefings

AcuGIS Cloud includes scheduling capabilities that automate report generation and distribution.

Scheduled reports may be:

- Generated automatically
- Exported in multiple formats
- Delivered to recipients
- Logged for auditing purposes

This functionality reduces manual effort while ensuring consistent reporting processes.

8.7 Report Distribution

Generating reports is only part of the reporting lifecycle.

Organizations must also distribute information to stakeholders.

Reporting workflows support:

- PDF delivery
- Spreadsheet delivery
- Email distribution
- Scheduled delivery
- Downloadable reports
- Embedded reports

This flexibility allows organizations to align reporting outputs with existing operational processes and user expectations.

The result is improved accessibility and reduced administrative overhead.

8.8 Integration with Spatial Data

A defining characteristic of AcuGIS Cloud reporting is direct integration with spatial information.

Reports can consume data originating from:

- GeoETL workflows
- GeoSync projects
- GeoVault datasets
- PostGIS environments
- Published services

This integration ensures that reporting activities remain connected to operational spatial information rather than relying on disconnected exports or duplicate data stores.

As datasets are updated, reporting workflows can reflect those changes without requiring additional synchronization procedures.

8.9 Map and Report Integration

Traditional GIS and reporting platforms frequently operate independently.

Maps provide spatial context while reports provide analytical summaries.

Users must often move between separate systems to obtain a complete understanding of operational conditions.

AcuGIS Cloud integrates these activities.

Published applications can expose reporting capabilities directly within the mapping environment.

Figure 13. Integrated mapping and reporting workflow

Interactive maps and operational reports can be presented within a unified user experience.

This integration provides several benefits:

Reduced Context Switching

Users remain within a single operational environment.

Improved Decision Support

Spatial and analytical information can be evaluated together.

Faster Access to Information

Reports become immediately accessible from relevant map features and workflows.

Better Operational Awareness

Maps and reports reinforce one another rather than competing for user attention.

8.10 Dashboard Capabilities

While reports provide structured outputs, dashboards provide continuous operational visibility.

Dashboards are particularly valuable when organizations require:

- Real-time monitoring
- Performance tracking
- Executive summaries
- Operational status indicators
- Trend analysis

Within AcuGIS Cloud, dashboards can combine:

- Maps
- Charts
- Tables
- Metrics
- Reports

into a unified operational workspace.

This approach allows users to monitor conditions continuously while retaining access to more detailed reporting workflows when necessary.

8.11 Spatial Analytics and Visualization

Operational reporting frequently benefits from visual representations of information.

Dashboard workflows may incorporate:

Maps

Spatial context and geographic distribution.

Charts

Comparisons, trends, and summaries.

Tables

Detailed operational records.

Metrics

Key performance indicators and operational status measurements.

By combining these elements, AcuGIS Cloud enables organizations to communicate information effectively to both technical and non-technical audiences.

8.12 Integration with Publishing

Reporting workflows are closely integrated with the publishing architecture.

Published applications can:

- Launch reports
- Display report outputs
- Pass parameters
- Present analytical results

This integration allows reporting to become part of the operational application rather than a separate destination.

Users can move seamlessly from map interaction to analytical output while remaining within the same workflow.

The result is a more cohesive user experience and improved operational efficiency.

8.13 Integration with APIs and Automation

Reports can also participate in automated workflows.

Common scenarios include:

- Scheduled execution
- Automated delivery
- External application integration
- Operational notifications
- Business process automation

These capabilities enable reporting activities to support larger organizational workflows while reducing manual effort.

As a result, reporting becomes an active component of operational systems rather than a passive output mechanism.

8.14 Operational Use Cases

Reporting and dashboard capabilities support a wide range of GIS applications.

Asset Management

Generate maintenance summaries, inspection reports, and performance metrics.

Environmental Monitoring

Track observations, trends, and compliance requirements.

Municipal Operations

Produce operational summaries, planning reports, and service analyses.

Community Data Portals

Provide public-facing analytical outputs alongside published maps.

Executive Reporting

Deliver concise summaries and key indicators to decision makers.

In each scenario, reporting serves as the mechanism through which operational spatial information is translated into actionable insight.

8.15 Architectural Benefits

The reporting architecture provides several important advantages.

Unified Data Sources

Reports consume the same operational information used by maps and services.

Reduced Duplication

Separate reporting databases and export workflows are minimized.

Integrated User Experience

Maps, reports, and dashboards function together.

Operational Automation

Scheduling and delivery workflows reduce manual effort.

Improved Decision Support

Users can combine spatial and analytical information within a single environment.

Together these capabilities strengthen the role of GIS as an operational decision-support platform.

8.16 Role Within AcuGIS Cloud

Reporting and dashboards represent the culmination of the AcuGIS Cloud lifecycle.

Data acquired through GeoETL, synchronized through GeoSync, managed through GeoVault, stored within PostGIS, and delivered through publishing services ultimately becomes available for analysis and decision support.

Reporting therefore serves as the mechanism through which operational information is transformed into organizational knowledge.

It completes the platform lifecycle by connecting spatial operations to business outcomes.

8.17 Summary

The reporting architecture of AcuGIS Cloud combines operational reporting, dashboard visualization, automation, and spatial integration within a unified platform.

Through Jasper-based reporting, dashboard capabilities, scheduling, distribution, and direct integration with operational datasets, AcuGIS Cloud enables organizations to move beyond map delivery and toward comprehensive decision support.

By treating reporting as an integral component of the spatial operations lifecycle, AcuGIS Cloud ensures that maps, services, reports, and analytics remain connected through a common operational foundation.

9. APIs, CLI, and Automation

9.1 Introduction

Modern GIS environments rarely operate in isolation.

Organizations increasingly require spatial systems to integrate with business applications, automation platforms, cloud services, analytics environments, and operational workflows.

As GIS becomes more deeply embedded within organizational processes, manual interaction alone is no longer sufficient.

Users require mechanisms for:

- System integration
- Process automation
- Application development
- Operational orchestration
- Scheduled execution
- Continuous workflows

AcuGIS Cloud addresses these requirements through a combination of APIs, command-line tools, and automation capabilities that extend platform functionality beyond the user interface.

These capabilities enable spatial operations to participate directly in broader enterprise workflows.

9.2 Automation as a Design Principle

Automation is a foundational principle of the AcuGIS Cloud architecture.

The platform is designed around the concept of deterministic operations that can be executed repeatedly and predictably.

Examples include:

- GeoETL pipelines
- GeoSync synchronization jobs

- Publishing workflows
- Report generation
- Dataset management operations

By exposing these capabilities through programmatic interfaces, organizations can automate routine activities while maintaining operational consistency.

Automation therefore serves as an extension of the deterministic workflow model introduced throughout the platform.

9.3 REST-Based APIs

AcuGIS Cloud provides API access to core platform capabilities.

APIs allow external applications to interact with platform services without requiring direct user interaction.

Typical API use cases include:

- Dataset discovery
- Data retrieval
- Workflow execution
- System integration
- Custom applications
- Operational dashboards

By exposing functionality through APIs, AcuGIS Cloud enables organizations to incorporate spatial operations into existing enterprise architectures.

Rather than requiring users to interact exclusively through browser interfaces, platform capabilities become accessible to a broader ecosystem of applications and services.

9.4 Dataset APIs

Managed datasets can be exposed through API-based access patterns.

These interfaces provide controlled access to operational information while maintaining governance and lifecycle controls.

Typical capabilities include:

- Dataset retrieval
- Feature access
- Metadata access
- Filtering
- Query operations
- Service integration

API-based access allows organizations to build custom applications while leveraging the same managed datasets used throughout the platform.

This approach reduces duplication and promotes consistency across operational systems.

9.5 Publishing APIs

Publishing workflows may also be integrated into automated processes.

Examples include:

- Project deployment
- Service publication
- Cache generation
- Update workflows
- Publication status monitoring

These capabilities enable organizations to automate portions of the publishing lifecycle while maintaining operational oversight.

As publishing environments grow, automation becomes increasingly valuable for maintaining consistency and reducing administrative effort.

9.6 Reporting APIs

Reporting workflows can be integrated into broader operational systems through programmatic interfaces.

Typical scenarios include:

- On-demand report generation
- Scheduled execution
- External application integration
- Dashboard integration
- Automated delivery

This flexibility allows reporting capabilities to participate directly in organizational business processes.

Rather than functioning solely as user-facing outputs, reports become reusable operational services.

9.7 The AcuGIS CLI

In addition to browser-based administration, AcuGIS Cloud provides command-line capabilities for operational users and administrators.

The AcuGIS CLI offers an alternative interaction model designed for:

- Administrators
- DevOps teams
- Automation workflows
- CI/CD environments
- Advanced operational users

Command-line interfaces are particularly valuable in environments where repeatability and automation are priorities.

Figure 14. AcuGIS command-line interface

The AcuGIS CLI provides administrative and operational access to platform services through a scriptable command-line environment.

The CLI extends platform capabilities beyond graphical interfaces and enables integration with existing operational tooling.

9.8 Administrative Operations

The CLI supports a variety of administrative workflows.

Typical examples include:

- Catalog inspection
- Dataset management
- Service administration
- Publishing operations
- Operational monitoring

These capabilities allow administrators to interact with platform resources efficiently while supporting repeatable operational procedures.

For organizations managing large deployments, command-line workflows can significantly reduce operational overhead.

9.9 Scriptable Workflows

Many GIS activities benefit from repeatable execution.

The AcuGIS CLI supports scripting and automation through standard command-line patterns.

Common examples include:

- Automated imports
- Publishing workflows
- Scheduled maintenance
- Monitoring operations
- Integration tasks

Because these workflows can be executed programmatically, organizations can incorporate spatial operations into existing automation frameworks.

This capability strengthens the platform's role within enterprise operational environments.

9.10 Deterministic Operations and Automation

A recurring theme throughout AcuGIS Cloud is deterministic execution.

Automation is most effective when operations behave predictably.

GeoETL pipelines, synchronization workflows, publishing activities, and reporting operations are designed to support repeatable execution.

This consistency enables organizations to automate workflows with confidence.

Rather than relying on undocumented manual procedures, automated processes operate against explicit definitions and managed resources.

The result is improved reliability and reduced operational risk.

9.11 Integration with Enterprise Systems

Organizations frequently need GIS platforms to interact with other operational systems.

Examples include:

- Asset management systems
- CRM platforms
- ERP environments
- Analytics platforms
- Data warehouses
- Custom applications

Through APIs and automation capabilities, AcuGIS Cloud can participate within these broader architectures.

This integration enables spatial information to become part of larger business workflows rather than remaining isolated within GIS-specific systems.

9.12 Scheduled Operations

Scheduling plays an important role in operational automation.

Recurring activities may include:

- Data synchronization
- Dataset imports

- Report generation
- Publishing updates
- Cache maintenance

AcuGIS Cloud supports scheduled execution of these activities, allowing organizations to automate repetitive processes while maintaining operational consistency.

By reducing reliance on manual execution, scheduling improves reliability and frees personnel to focus on higher-value activities.

9.13 CI/CD and DevOps Workflows

As GIS systems increasingly adopt modern software development practices, integration with DevOps workflows becomes more important.

The CLI and API architecture support:

- Automated deployments
- Configuration management
- Environment promotion
- Testing workflows
- Operational monitoring

These capabilities allow spatial operations to participate in the same lifecycle management practices used throughout modern software environments.

This alignment simplifies administration while improving operational maturity.

9.14 Operational Benefits

The combination of APIs, command-line tools, and automation provides several advantages.

Repeatability

Operations can be executed consistently across environments.

Integration

Platform capabilities can be incorporated into existing systems and workflows.

Scalability

Automation reduces administrative effort as deployments grow.

Reliability

Deterministic operations improve predictability and reduce operational risk.

Flexibility

Organizations can interact with the platform through user interfaces, APIs, command-line tools, or automated workflows.

Together these capabilities support a wide variety of operational requirements while preserving consistency across the platform.

9.15 Role Within AcuGIS Cloud

APIs, CLI capabilities, and automation are not separate products.

They are enabling technologies that extend every major component of the platform.

They support:

- GeoETL
- GeoSync
- GeoVault
- PostGIS Service
- Publishing services
- Reporting workflows

By providing multiple interaction models, AcuGIS Cloud allows organizations to choose the operational approach that best fits their requirements.

This flexibility is essential for environments ranging from small operational teams to enterprise-scale deployments.

9.16 Summary

The API, CLI, and automation capabilities of AcuGIS Cloud extend the platform beyond traditional browser-based GIS administration.

Through REST-based interfaces, command-line tooling, scheduling, and integration capabilities, organizations can incorporate spatial operations directly into enterprise workflows and automated processes.

These capabilities reinforce the platform's deterministic architecture while enabling scalable, repeatable, and interoperable operational GIS environments.

As a result, AcuGIS Cloud functions not only as a collection of GIS tools, but as a programmable spatial operations platform capable of participating in broader organizational systems and processes.

10. Deployment Architecture

10.1 Introduction

Modern GIS deployments must balance flexibility, scalability, security, and operational control.

Organizations increasingly require platforms that can support a variety of deployment models while maintaining a consistent operational experience. Some environments prioritize self-hosting and infrastructure control, while others emphasize simplified administration, cloud-based deployment, or hybrid architectures.

AcuGIS Cloud is designed to support these diverse requirements through a modular, service-oriented deployment architecture.

Rather than requiring a single deployment model, the platform allows organizations to adopt infrastructure strategies that align with operational, security, and governance requirements.

This flexibility enables AcuGIS Cloud to support deployments ranging from small departmental environments to enterprise-scale spatial operations platforms.

10.2 Architectural Principles

The deployment architecture is guided by several core principles.

Self-Hosted First

Organizations should retain control over their infrastructure, data, and operational processes.

While managed environments may be supported, the platform is designed so that organizations can deploy and operate AcuGIS Cloud within their own infrastructure when required.

Modular Services

Platform capabilities are organized into independent services that can evolve and scale according to operational requirements.

Open Infrastructure

The platform is built upon widely adopted open-source technologies including:

- PostgreSQL
- PostGIS
- QGIS
- GeoServer
- Apache
- Linux

This approach reduces dependency on proprietary infrastructure while simplifying integration with existing environments.

Operational Independence

Individual services can be deployed, updated, monitored, and managed independently while remaining integrated through platform workflows.

10.3 Service-Oriented Architecture

AcuGIS Cloud follows a service-oriented architecture in which major platform components operate as distinct services.

Typical services include:

- GeoETL
- GeoSync
- GeoVault
- PostGIS Service
- QGIS Publisher
- GeoServer
- Reporting Services
- Dashboard Services

These services communicate through APIs, shared databases, and operational workflows.

This architecture provides several benefits:

- Improved maintainability
- Independent scaling
- Simplified upgrades
- Operational flexibility
- Clear separation of responsibilities

By decoupling major functions, the platform can evolve without requiring changes across every component.

10.4 Reference Deployment

A typical AcuGIS Cloud deployment consists of several integrated layers.

Presentation Layer

User-facing services including:

- GeoVault
- QGIS Publisher
- Dashboards
- Administrative interfaces

Application Layer

Operational services including:

- GeoETL
- GeoSync
- Reporting Services
- Publishing Services

Data Layer

Persistent storage including:

- PostgreSQL
- PostGIS
- Dataset repositories
- Operational metadata

Service Layer

Publication and delivery services including:

- QGIS Server
- GeoServer
- Tile services
- APIs

Figure 5. Deployment Architecture

Reference deployment architecture showing the relationship between user-facing applications, operational services, publishing infrastructure, and PostGIS environments.

This layered architecture supports both operational simplicity and long-term scalability.

10.5 Linux Deployment Model

AcuGIS Cloud is designed primarily for Linux-based environments.

Typical deployments include:

- Ubuntu LTS
- Rocky Linux
- Enterprise Linux distributions

Linux provides a stable and widely adopted foundation for spatial infrastructure and supports the open-source technologies used throughout the platform.

This deployment model allows organizations to leverage existing operational expertise while maintaining flexibility in infrastructure selection.

10.6 Containerized Deployments

Many organizations prefer containerized infrastructure for deployment and operations.

AcuGIS Cloud components can be deployed using container technologies where operational requirements support this approach.

Benefits may include:

- Simplified installation
- Consistent environments
- Improved portability
- Deployment automation
- Environment isolation

Containerized deployments can be particularly valuable in cloud-native environments and automated infrastructure workflows.

10.7 Traditional Deployments

Not all organizations require containerized infrastructure.

AcuGIS Cloud also supports traditional deployment models in which services are installed directly on operating systems and managed through standard administrative procedures.

Advantages of traditional deployments may include:

- Simpler troubleshooting
- Direct access to services
- Familiar administration workflows
- Reduced abstraction

This flexibility allows organizations to select deployment approaches that align with operational preferences and existing infrastructure standards.

10.8 Database Architecture

PostGIS serves as the operational database foundation of the platform.

Deployment architectures may include:

- Single PostGIS environments
- Multiple database environments
- Development and production separation
- Department-specific databases
- Dedicated reporting databases

PostGIS Service provides management capabilities that support these deployment patterns while maintaining integration with platform services.

This architecture enables organizations to scale database infrastructure according to operational requirements.

10.9 Publishing Infrastructure

Publishing services can be deployed independently from administrative and processing services.

This separation allows organizations to optimize infrastructure for specific workloads.

Examples include:

Publishing Servers

Focused on map delivery and service performance.

Processing Servers

Dedicated to GeoETL and synchronization workflows.

Reporting Servers

Optimized for report generation and distribution.

Administrative Servers

Focused on management and operational interfaces.

Separating workloads can improve scalability and operational resilience.

10.10 Security and Access Control

Security considerations are integrated throughout the deployment architecture.

Typical security capabilities include:

- Authentication
- Authorization
- Role-based access control
- Database security
- Service isolation
- Administrative controls

These capabilities help organizations maintain governance while supporting operational access to spatial information.

Security models can be adapted to organizational requirements without fundamentally altering platform architecture.

10.11 Network Architecture

Deployment architectures may support a variety of network topologies.

Examples include:

Internal Deployments

Accessible only within organizational networks.

Public Deployments

Providing public-facing maps, services, and reports.

Hybrid Deployments

Combining internal operational services with externally accessible publishing infrastructure.

Multi-Tier Deployments

Separating presentation, application, and data layers across network boundaries.

This flexibility allows organizations to align deployment architectures with operational and security requirements.

10.12 Scalability Considerations

As spatial operations grow, deployment architectures must support increasing workloads.

Scalability requirements may include:

- Larger datasets
- More users
- Increased synchronization activity
- Additional reports
- Expanded publishing infrastructure

The modular architecture of AcuGIS Cloud allows services to scale independently according to operational demand.

For example:

- Publishing infrastructure can be expanded without modifying GeoETL workflows.
- Reporting services can scale independently of synchronization services.
- Database resources can be increased without affecting user interfaces.

This approach supports long-term growth while maintaining operational flexibility.

10.13 High Availability and Resilience

Organizations with mission-critical GIS requirements may require enhanced availability.

The deployment architecture supports common resilience strategies including:

- Database backups
- Service redundancy
- Infrastructure monitoring
- Disaster recovery procedures
- Operational failover strategies

Specific implementations vary according to organizational requirements and infrastructure standards.

The modular nature of the platform simplifies the application of these strategies across individual services.

10.14 Monitoring and Operations

Operational visibility is essential for managing spatial infrastructure.

Typical monitoring activities include:

- Service health monitoring
- Database monitoring
- Synchronization monitoring
- Report execution monitoring
- Publishing status monitoring

Monitoring capabilities enable organizations to identify issues quickly and maintain operational continuity.

As deployments grow, these capabilities become increasingly important for maintaining reliability.

10.15 Cloud and Hybrid Deployments

Although AcuGIS Cloud supports self-hosted deployments, organizations may also choose managed or hybrid deployment models.

Examples include:

- Private cloud infrastructure
- Public cloud environments
- Hybrid deployments
- Managed hosting services

The platform architecture remains consistent across these deployment models.

This consistency simplifies migration and reduces operational complexity when infrastructure strategies evolve.

10.16 Architectural Benefits

The deployment architecture provides several key advantages.

Flexibility

Organizations can select deployment models that align with operational requirements.

Scalability

Services can scale independently.

Maintainability

Modular architecture simplifies upgrades and administration.

Interoperability

Open standards and open-source technologies reduce vendor lock-in.

Operational Control

Organizations retain control over infrastructure and deployment decisions.

Together these capabilities support a wide variety of GIS deployment scenarios.

10.17 Role Within AcuGIS Cloud

The deployment architecture provides the operational foundation upon which all platform capabilities are built.

GeoETL, GeoSync, GeoVault, PostGIS Service, publishing infrastructure, reporting systems, APIs, and automation workflows all depend upon a deployment architecture that supports reliability, scalability, and operational control.

By combining modular services with flexible deployment models, AcuGIS Cloud enables organizations to implement spatial operations platforms that align with both current requirements and future growth.

10.18 Summary

The deployment architecture of AcuGIS Cloud is designed to support self-hosted, cloud-based, and hybrid GIS environments through a modular service-oriented approach.

By leveraging open technologies, flexible deployment models, and clear separation of responsibilities, the platform provides a scalable foundation for spatial operations across a wide range of organizational environments.

This architecture enables organizations to deploy, manage, and evolve GIS infrastructure while maintaining operational continuity, governance, and long-term flexibility.

11. Example Operational Workflows

11.1 Introduction

The preceding sections described the architecture, components, and deployment model of AcuGIS Cloud.

While understanding individual platform capabilities is important, the true value of the platform emerges when those capabilities are combined into operational workflows.

This section presents representative workflows demonstrating how organizations can acquire, synchronize, transform, manage, publish, and report on spatial information using AcuGIS Cloud.

These examples illustrate practical usage patterns rather than exhaustive implementation details.

Together they demonstrate how the platform supports the complete lifecycle of operational spatial data.

11.2 Open Data Acquisition and Publishing

One of the most common GIS activities involves acquiring external datasets and making them available through operational applications.

Examples include:

- Municipal open data portals
- Census datasets
- Environmental monitoring data
- Transportation information

- Community datasets

In many organizations this process involves multiple disconnected systems and extensive manual effort.

AcuGIS Cloud consolidates these activities into a unified workflow.

Workflow

1. Dataset imported through GeoVault.
2. Optional transformation executed through GeoETL.
3. Dataset stored and versioned within GeoVault.
4. Data loaded into PostGIS.
5. Dataset published through QGIS Publisher or GeoServer.
6. Published application made available to users.

Benefits

- Reduced manual processing
- Version-controlled datasets
- Repeatable workflows
- Faster publication cycles
- Consistent operational procedures

Figure 15. Open data acquisition and publishing workflow

Figure 1 — Platform Overview

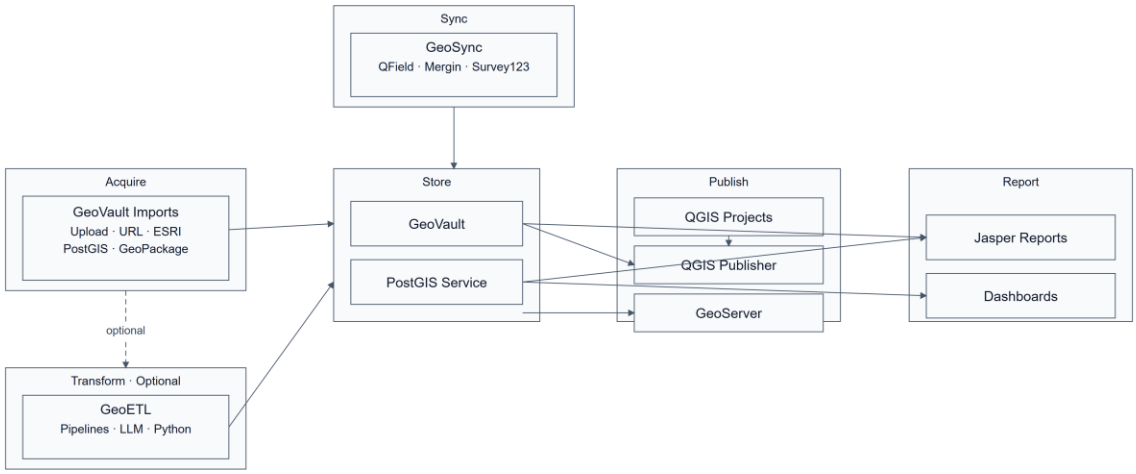

External datasets can be imported, transformed, managed, and published through a unified operational workflow.

11.3 Field Data Collection Workflow

Many organizations rely on mobile field collection systems to gather operational information.

Examples include:

- Asset inspections
- Environmental surveys
- Infrastructure inventories
- Community reporting
- Utility management

Historically, synchronizing field data with enterprise GIS environments has required custom scripts and manual workflows.

GeoSync provides a unified synchronization architecture for operational field data.

Workflow

1. Field crews collect information using mobile applications.
2. Changes are synchronized through GeoSync.
3. GeoPackage differences are processed.
4. Updates are applied to PostGIS.
5. Published maps automatically reflect changes.
6. Dashboards and reports consume updated information.

Supported Sources

- QField
- Mergin Maps
- Survey123

Benefits

- Automated synchronization
- Reduced manual imports
- Consistent operational data
- Faster availability of updates
- Improved field-to-office workflows

Figure 16. Field synchronization workflow

Figure A3 — GeoSync Workflow

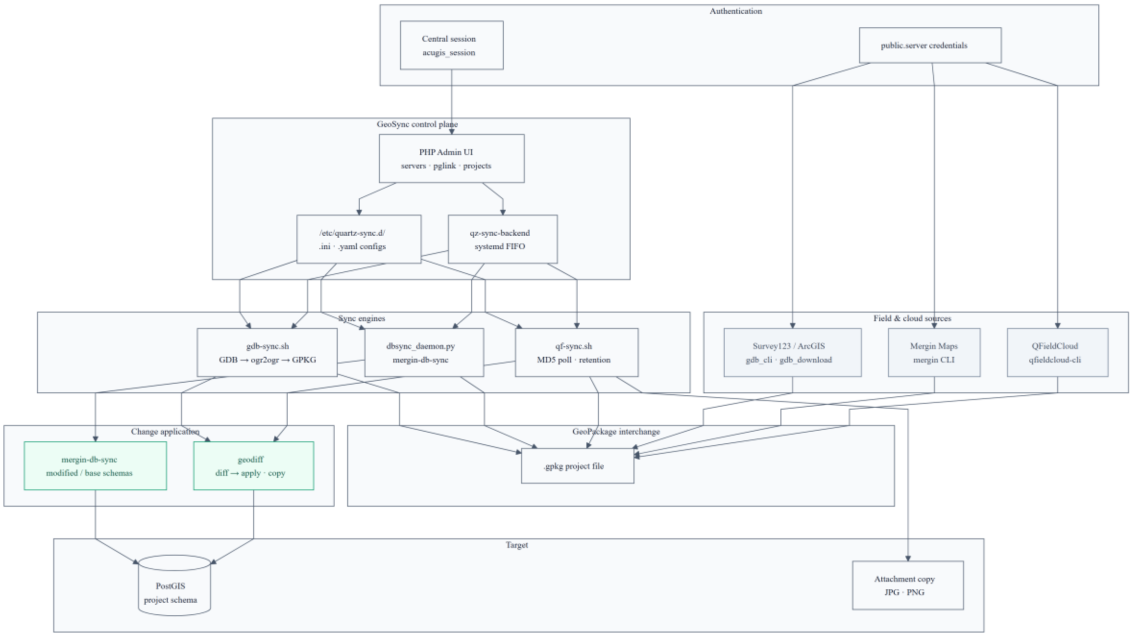

GeoSync automates the movement of operational field information into enterprise spatial databases.

11.4 Dataset Transformation Workflow

Organizations frequently receive datasets that require modification before operational use.

Common examples include:

- Schema normalization
- Field renaming
- Attribute calculations
- Data filtering
- Geometry cleanup
- Dataset enrichment

GeoETL provides deterministic transformation workflows that standardize these activities.

Workflow

1. Dataset imported into GeoVault.
2. Dataset inspected within GeoETL.
3. Transformations defined using deterministic operations.
4. Optional LLM-assisted interpretation generates pipeline definitions.
5. Preview validates proposed changes.
6. Approved transformations execute.
7. Resulting dataset stored and versioned.

Example Transformations

- Rename fields
- Add calculated attributes
- Remove columns
- Filter records
- Dissolve features
- Reproject geometries

Benefits

- Repeatable execution
- Transformation transparency
- Validation before import
- Reduced manual editing
- Reusable pipelines

Figure 17. Deterministic transformation workflow

Figure A2 — GeoETL Workflow

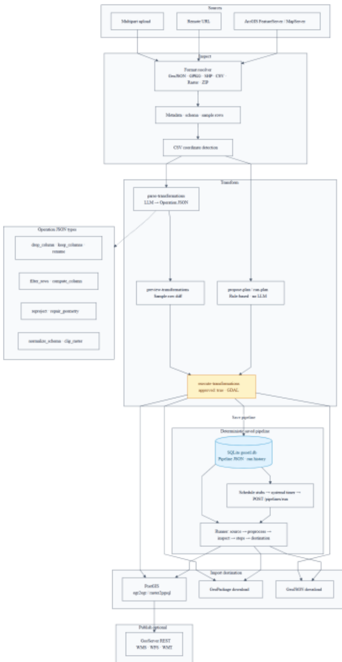

GeoETL enables controlled transformation of datasets prior to operational deployment.

11.5 Operational Editing Workflow

Many datasets continue to evolve after initial publication.

Organizations require mechanisms for updating information while preserving data lineage and operational integrity.

GeoVault provides version management capabilities that support ongoing dataset maintenance.

Workflow

1. Dataset imported into GeoVault.
2. Dataset published for operational use.
3. Users perform updates and corrections.
4. New versions are created.
5. Previous versions remain available.
6. Rollback operations can restore earlier states.
7. Published services continue to reference managed datasets.

Benefits

- Dataset history
- Controlled updates
- Operational continuity
- Rollback capability
- Improved governance

This workflow is particularly valuable in environments where datasets are updated frequently and operational reliability is critical.

Figure 18. Dataset versioning and operational editing

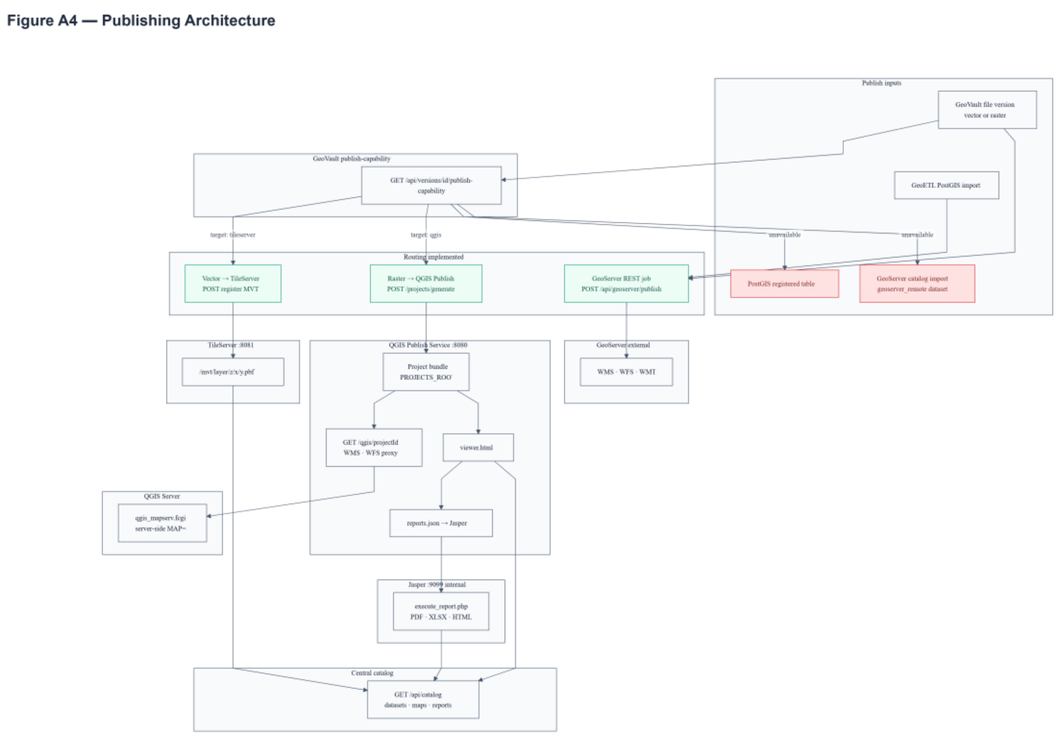

GeoVault supports controlled evolution of operational datasets through managed version histories.

11.6 QGIS Publishing Workflow

Many organizations maintain existing QGIS-based cartographic and analytical workflows.

AcuGIS Cloud extends these workflows by enabling direct publication of QGIS projects.

Workflow

1. QGIS project prepared in QGIS Desktop.
2. Project uploaded through QGIS Publisher.
3. Project published as an operational application.
4. Services exposed through standard interfaces.
5. End users access published maps through browsers and applications.

Published Outputs

- Interactive map applications
- WMS services
- WFS services
- Operational dashboards
- Embedded applications

Benefits

- Reuse of existing QGIS investments
- Rapid deployment
- Reduced duplication
- Consistent cartography
- Integrated publication workflow

Figure 19. QGIS publication workflow

Figure A5 — Deployment Architecture

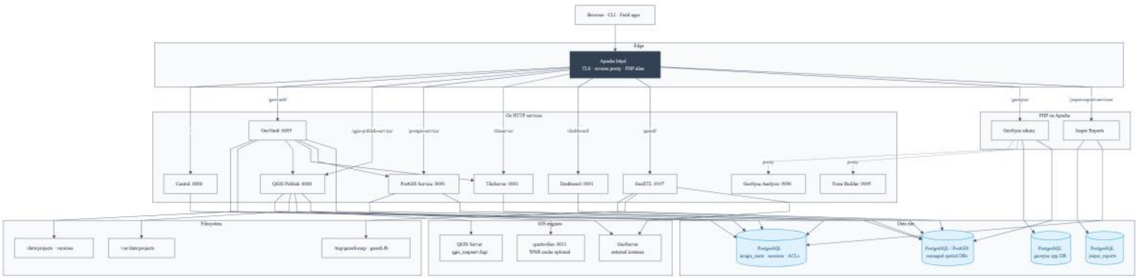

QGIS Publisher enables organizations to transform desktop GIS projects into operational applications.

11.7 Reporting Workflow

Operational reporting transforms spatial information into organizational insight.

Within AcuGIS Cloud, reporting workflows operate directly against managed operational datasets.

Workflow

1. Dataset stored in PostGIS.
2. Report template published.
3. Parameters supplied by users or schedules.
4. Report executed.
5. Results delivered as PDF, spreadsheet, or embedded output.
6. Reports distributed to stakeholders.

Common Outputs

- Operational summaries
- Inspection reports
- Asset inventories
- Environmental reports
- Compliance reporting
- Executive summaries

Benefits

- Reduced duplication
- Direct access to operational data
- Automated generation
- Consistent reporting

Figure 20. Operational reporting workflow

Reporting capabilities provide analytical outputs directly from managed operational datasets.

11.8 Dashboard Monitoring Workflow

Operational users frequently require continuous visibility into current conditions.

Dashboards provide an alternative to scheduled reports by presenting information in real time.

Workflow

1. Operational datasets updated through imports or synchronization.
2. Data stored within PostGIS.
3. Dashboard components query operational information.
4. Maps, charts, tables, and metrics update.
5. Users monitor conditions through a unified interface.

Dashboard Components

- Maps
- Charts

- Tables
- Metrics
- Reports

Benefits

- Continuous visibility
- Operational awareness
- Reduced reporting latency
- Improved decision support

This workflow is particularly useful for monitoring active operations and rapidly changing conditions.

Figure 21. Dashboard monitoring workflow

Dashboards provide continuous access to operational spatial information.

11.9 End-to-End Spatial Operations Workflow

The most comprehensive example combines all major platform components into a single operational lifecycle.

Scenario

A municipality wishes to collect field information, update operational datasets, publish applications, and generate reports.

Workflow

Step 1 — Acquire

Data is imported through GeoVault from external sources.

Step 2 — Sync

Field observations are synchronized through GeoSync.

Step 3 — Transform

GeoETL standardizes and validates incoming information.

Step 4 — Store

GeoVault and PostGIS maintain operational datasets.

Step 5 — Publish

QGIS Publisher and GeoServer expose maps and services.

Step 6 — Report

Dashboards and Jasper Reports deliver analytical outputs.

This workflow mirrors the platform lifecycle described throughout this white paper:

Acquire → Sync → Transform → Store → Publish → Report

The result is a unified spatial operations platform that supports the complete lifecycle of operational information.

Figure 22. End-to-end spatial operations workflow

The AcuGIS Cloud lifecycle integrates acquisition, synchronization, transformation, storage, publishing, and reporting within a single operational environment.

11.10 Architectural Significance

The workflows presented in this section illustrate an important characteristic of AcuGIS Cloud.

The platform is not simply a collection of independent GIS tools.

Instead, its components are designed to operate together as part of a cohesive operational architecture.

This integration provides several advantages:

Reduced Complexity

Fewer disconnected systems and manual procedures.

Improved Consistency

Common operational patterns across workflows.

Faster Delivery

Data moves more rapidly from acquisition to decision support.

Better Governance

Versioning, management, and operational controls remain integrated throughout the lifecycle.

Enhanced Decision Support

Maps, reports, dashboards, and applications operate against a shared operational foundation.

11.11 Summary

The workflows presented in this section demonstrate how AcuGIS Cloud supports practical operational GIS activities across the complete spatial data lifecycle.

Whether importing open data, synchronizing field observations, transforming datasets, publishing applications, or generating reports, organizations can leverage a common operational architecture that promotes consistency, repeatability, and efficiency.

These workflows illustrate the central goal of AcuGIS Cloud: enabling deterministic spatial operations through an integrated platform that connects data acquisition, synchronization, transformation, storage, publication, and decision support.

Conclusion

Deterministic GIS Pipelines and Spatial Operations

Geographic information systems have evolved significantly over the past several decades.

Organizations today are expected to manage increasing volumes of spatial information, support mobile field operations, integrate diverse data sources, publish applications to a broad range of users, and provide timely analytical outputs that support operational and strategic decision-making.

While the technologies used to support these activities have matured, many organizations continue to rely on fragmented workflows composed of disconnected tools, manual processes, custom scripts, and isolated data stores.

These approaches often introduce unnecessary complexity, increase operational risk, and make it difficult to maintain consistency across the spatial data lifecycle.

AcuGIS Cloud was developed to address these challenges.

Rather than focusing on a single aspect of GIS operations, the platform provides an integrated framework that supports the complete lifecycle of operational spatial information.

From acquisition and synchronization through transformation, storage, publication, and reporting, AcuGIS Cloud enables organizations to manage spatial operations through a unified architecture.

A Unified Spatial Operations Platform

Throughout this white paper, AcuGIS Cloud has been presented not as a collection of independent applications, but as a cohesive operational platform.

Each major component serves a distinct purpose within the broader spatial lifecycle:

GeoVault

Provides acquisition, management, versioning, and governance for operational datasets.

GeoETL

Supports deterministic transformation workflows through reusable pipelines, validation, and optional LLM-assisted configuration.

GeoSync

Connects field data collection environments with enterprise spatial databases through automated synchronization workflows.

PostGIS Service

Provides a managed operational database foundation for spatial information.

QGIS Publisher and GeoServer

Enable publication of maps, services, and operational applications.

Jasper Reports and Dashboards

Transform operational data into analytical outputs and decision-support tools.

Together these components form a unified environment for managing operational spatial information.

Deterministic Operations as a Core Principle

A recurring theme throughout this white paper has been determinism.

Many GIS workflows rely heavily on manual procedures that are difficult to reproduce, validate, or automate.

AcuGIS Cloud takes a different approach.

Transformation pipelines, synchronization activities, publication workflows, and reporting operations are designed to be explicit, repeatable, and transparent.

This emphasis on deterministic execution provides several benefits:

- Improved reliability
- Greater operational consistency
- Enhanced auditability
- Reduced administrative effort
- Easier automation

By treating spatial operations as managed processes rather than ad hoc activities, organizations can improve both efficiency and governance.

Bridging Traditional GIS and Operational Systems

Historically, GIS platforms have often functioned as specialized environments separate from broader organizational workflows.

Data collection, publication, reporting, and analysis frequently occur in different systems managed by different teams.

AcuGIS Cloud seeks to bridge these gaps.

Through APIs, automation capabilities, reporting integration, synchronization workflows, and publishing services, the platform enables spatial information to participate directly in operational processes.

This integration allows GIS to move beyond visualization and become an active component of organizational decision-making and service delivery.

Open Architecture and Flexibility

The platform is built upon widely adopted open technologies including:

- PostgreSQL
- PostGIS
- QGIS
- GeoServer
- Apache
- Linux

This open architecture provides flexibility while reducing dependency on proprietary technologies.

Organizations can deploy the platform within a variety of infrastructure environments, integrate with existing systems, and adapt workflows to meet specific operational requirements.

By leveraging open standards and open-source technologies, AcuGIS Cloud promotes interoperability, transparency, and long-term sustainability.

Supporting the Complete Lifecycle

A central objective of the platform is support for the complete lifecycle of operational spatial information.

This lifecycle can be summarized as:

Acquire → Sync → Transform → Store → Publish → Report

Data enters the platform through imports, synchronization services, and operational workflows.

It is transformed, validated, and managed through deterministic processes.

It is stored within managed spatial environments.

It is published through maps, services, and applications.

Finally, it is transformed into actionable information through reports, dashboards, and analytical outputs.

This lifecycle provides a framework through which organizations can standardize and scale spatial operations.

From Data Management to Decision Support

The ultimate purpose of spatial information is not storage or publication alone.

Its value emerges when it supports operational awareness, planning, analysis, and decision-making.

AcuGIS Cloud is designed with this objective in mind.

By connecting acquisition, synchronization, transformation, publication, and reporting within a single architecture, the platform helps organizations move more efficiently from raw data to actionable insight.

Maps, reports, dashboards, and applications operate against a shared operational foundation, ensuring consistency across the information lifecycle.

The result is improved visibility, better governance, and more effective use of spatial information.

Looking Forward

The demand for operational spatial intelligence continues to grow across both public and private sectors.

Organizations increasingly require platforms that support:

- Field operations
- Data integration
- Workflow automation
- Spatial analytics
- Operational reporting
- Enterprise interoperability

Meeting these requirements requires more than individual GIS tools.

It requires an architecture capable of supporting the complete spatial operations lifecycle.

AcuGIS Cloud represents one approach to this challenge.

By combining deterministic pipelines, operational synchronization, managed spatial storage, publishing infrastructure, and integrated reporting capabilities, the platform provides a foundation for modern spatial operations.

Final Remarks

AcuGIS Cloud was designed around a simple idea:

Spatial operations should be repeatable, manageable, and integrated.

Through deterministic workflows, open architecture, and lifecycle-oriented design, the platform enables organizations to move beyond isolated GIS activities and toward a more operational approach to spatial information management.

The result is a platform that connects data acquisition, synchronization, transformation, storage, publication, and decision support within a single operational framework.

As organizations continue to expand their use of spatial information, platforms that unify these capabilities will play an increasingly important role in delivering reliable, scalable, and actionable spatial intelligence.

AcuGIS Cloud

Deterministic GIS Pipelines and Spatial Operations

Figure 1 — GeoETL Deterministic Transformation Design

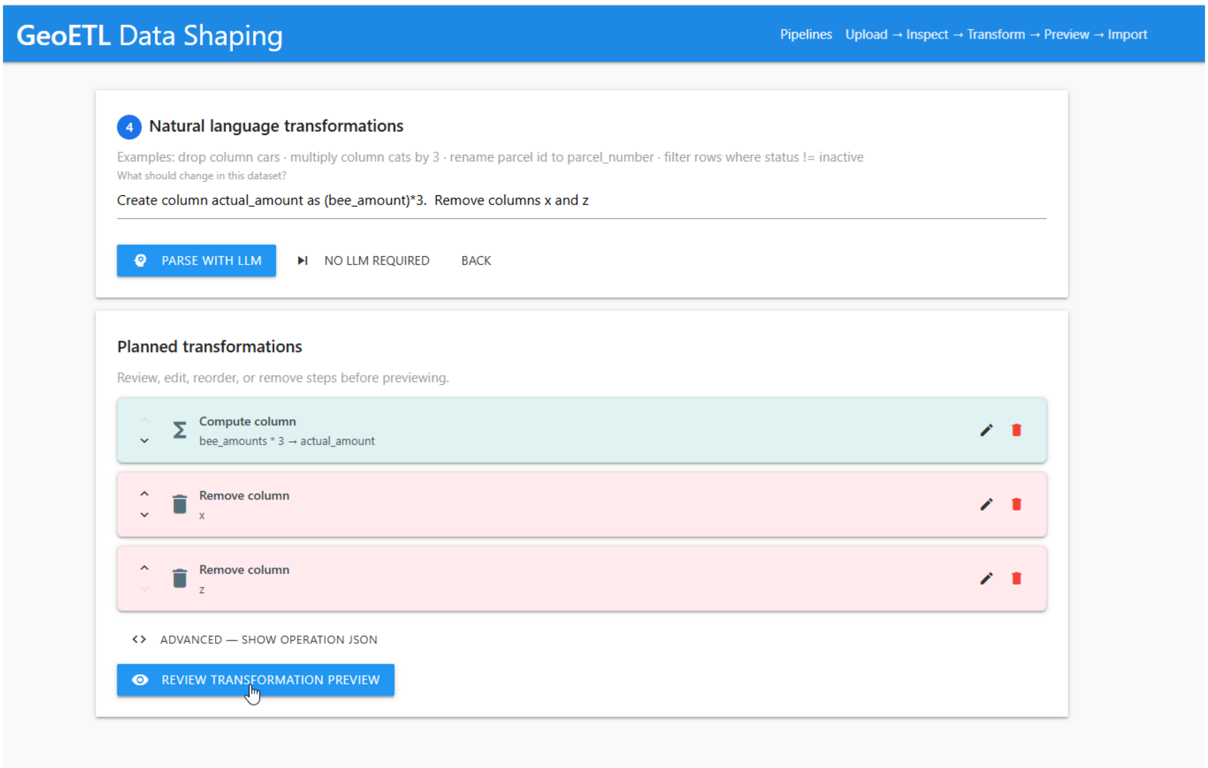

GeoETL converts transformation requests into explicit operations that can be reviewed, edited, validated, and executed.

Figure 2 — GeoETL Transformation Preview

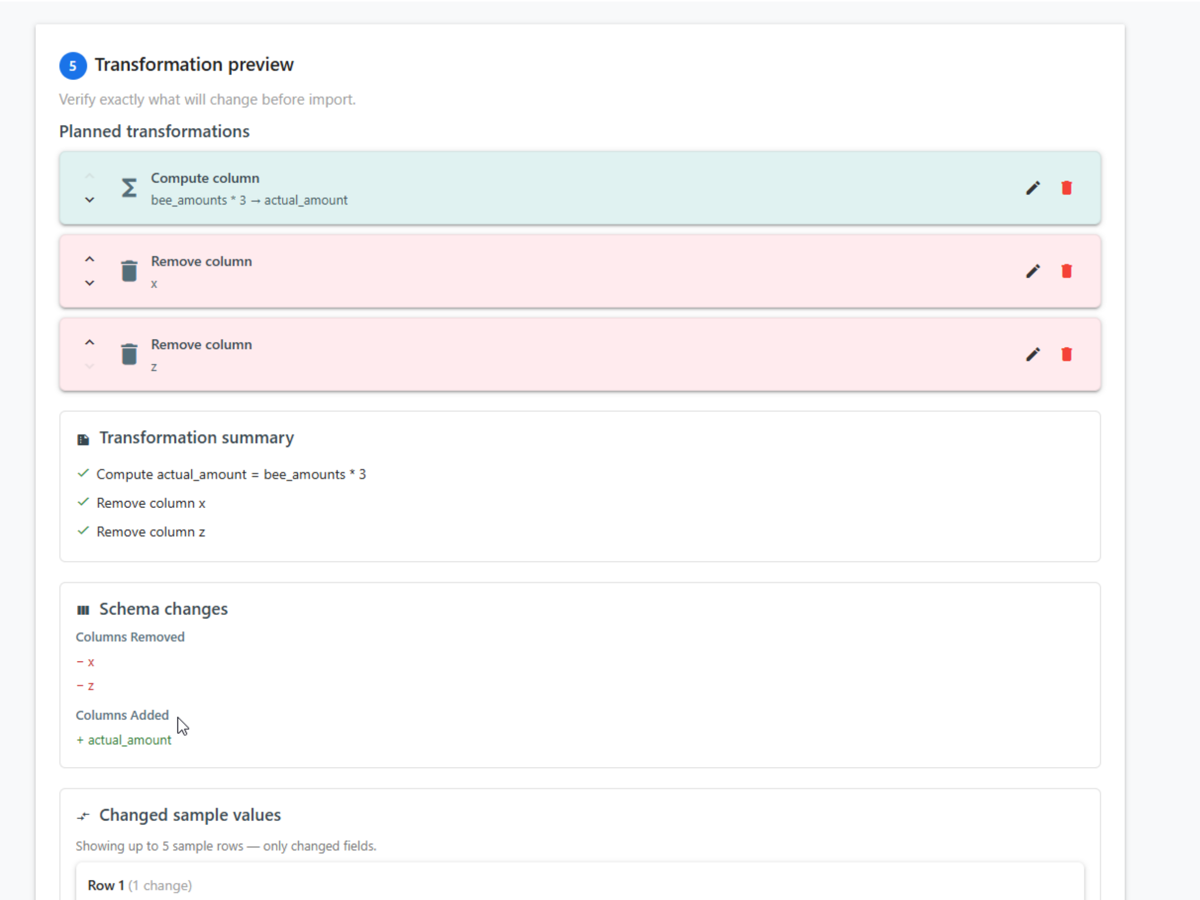

Transformation preview displays schema modifications, calculated fields, and sample changes prior to execution.

Figure 3 — GeoETL Import Execution

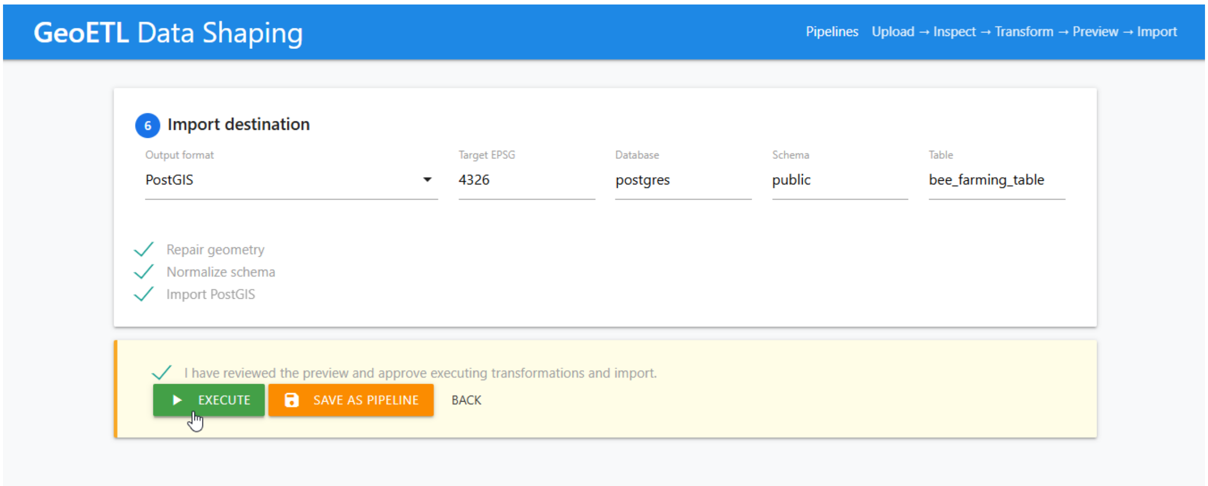

Approved transformations may be executed immediately or saved as reusable pipelines.

Figure 4 — GeoSync Project Wizard

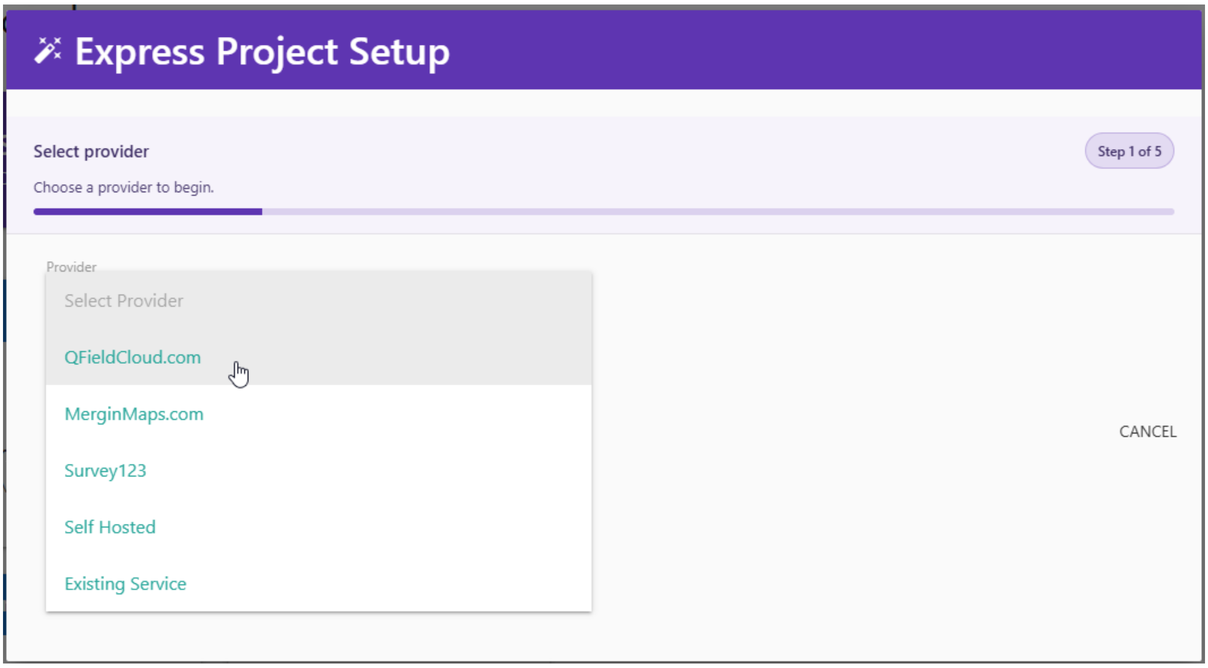

GeoSync supports synchronization workflows for QFieldCloud, Mergin Maps, Survey123, and self-hosted environments.

Figure 5 — GeoSync Operational Configuration

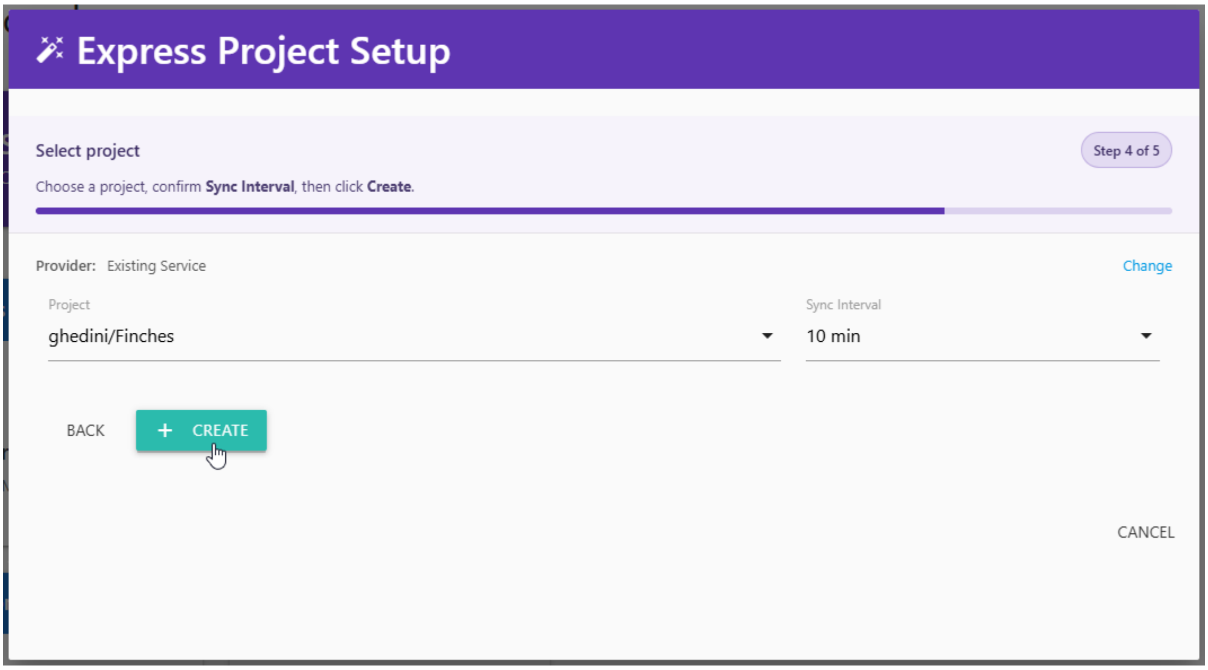

Synchronization schedules are configured at the project level and executed automatically.

Figure 6 — GeoVault Dataset Catalog

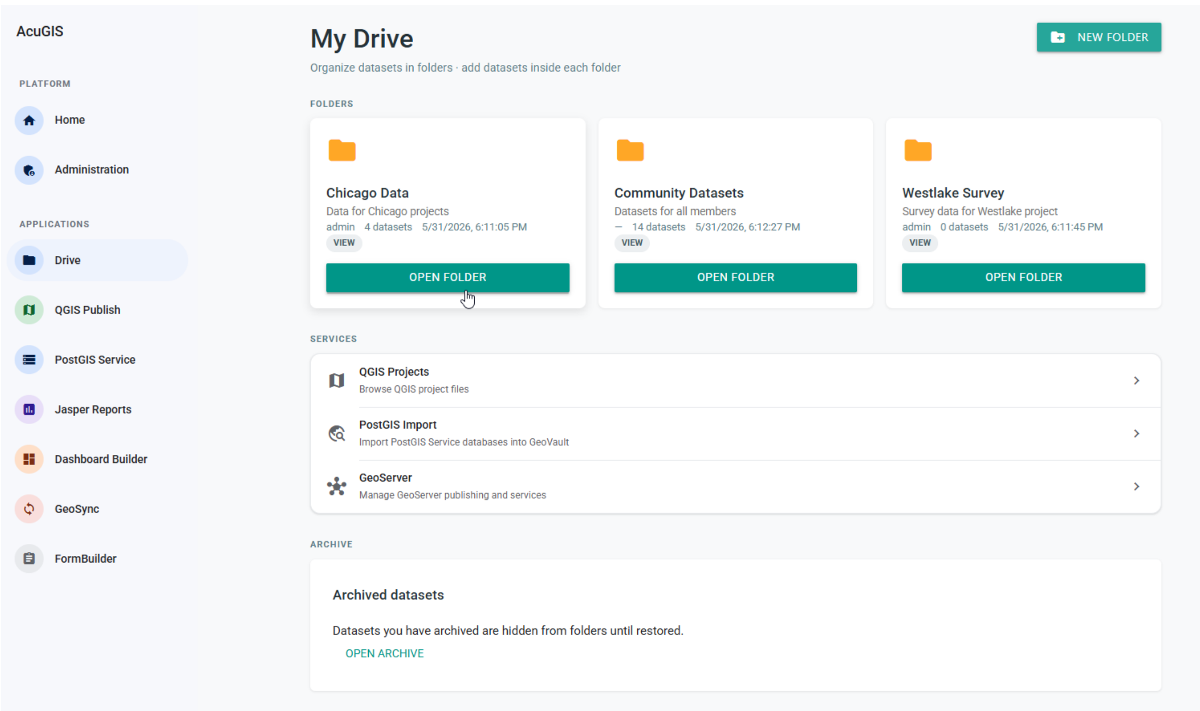

GeoVault organizes datasets into operational folders while exposing publishing and database services through a unified interface.

Figure 7 — GeoVault Version Management

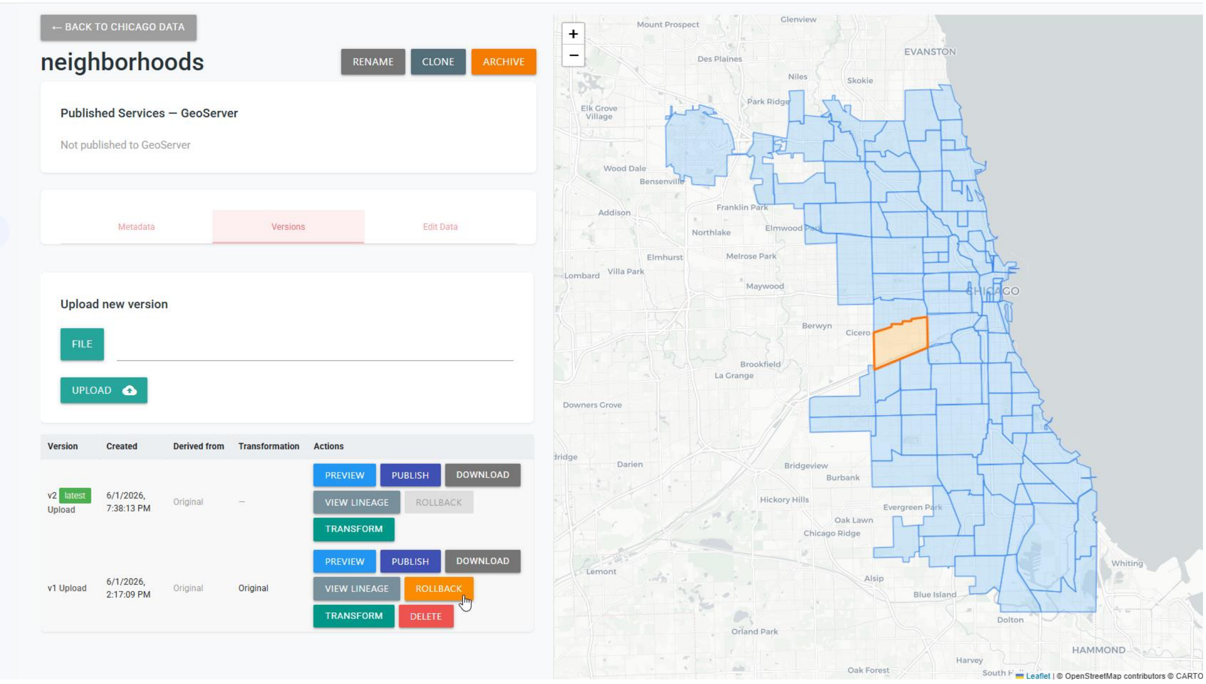

Dataset versions preserve lineage, support rollback, and maintain transformation history.

Figure 8 — GeoVault Spatial Editing

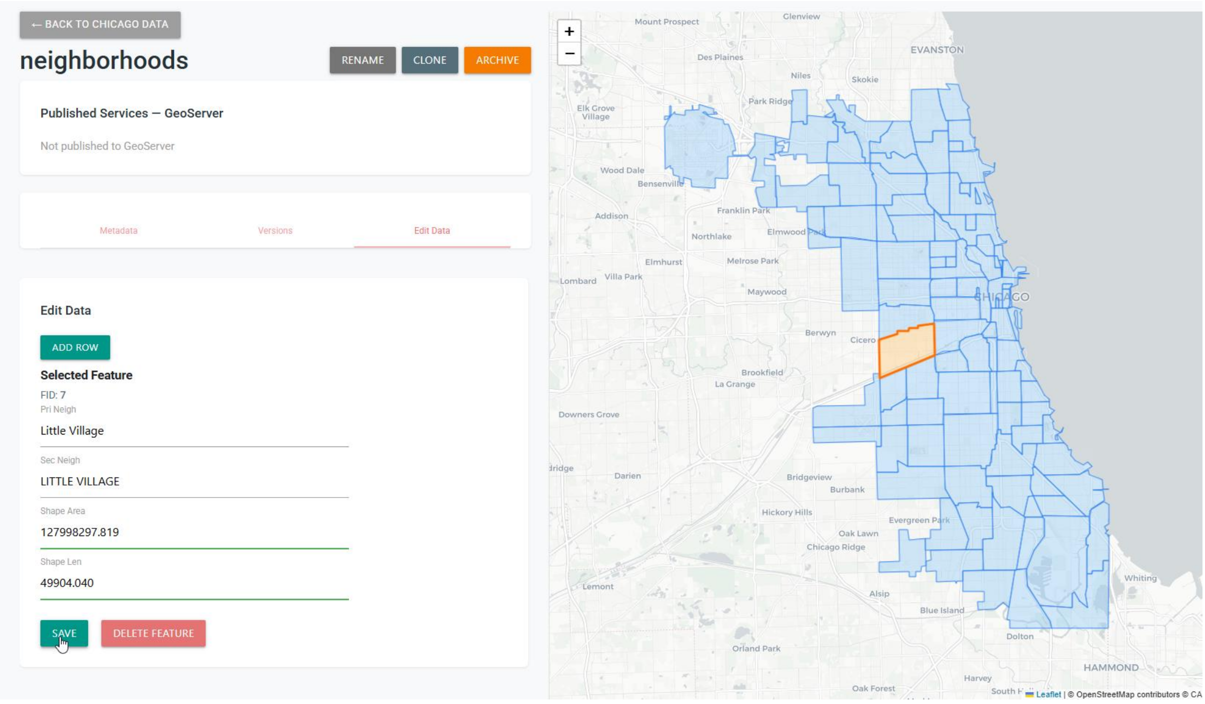

GeoVault provides browser-based editing of spatial features and associated attributes.

Figure 9 — QGIS Publisher

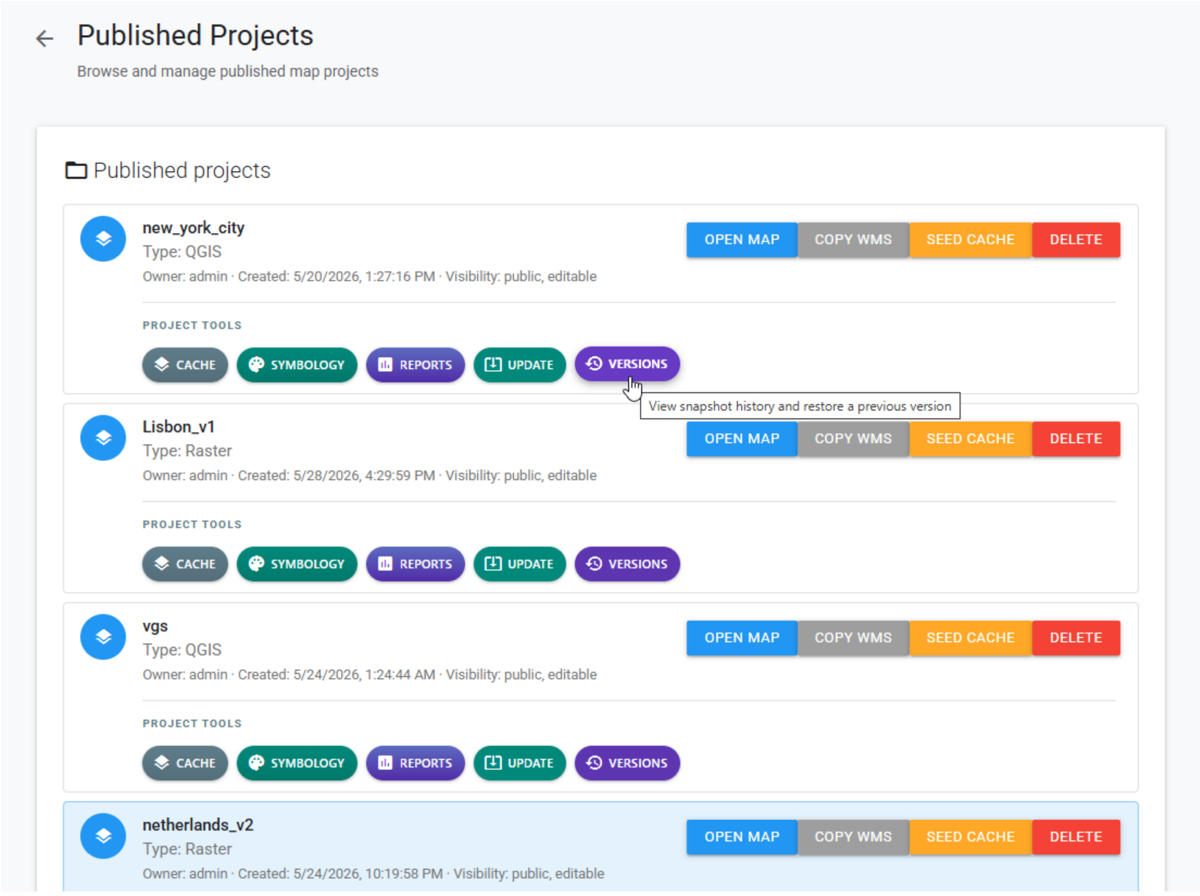

Published projects may be updated, versioned, cached, and integrated with reporting workflows.

Figure 10 — Published Map + Report Integration

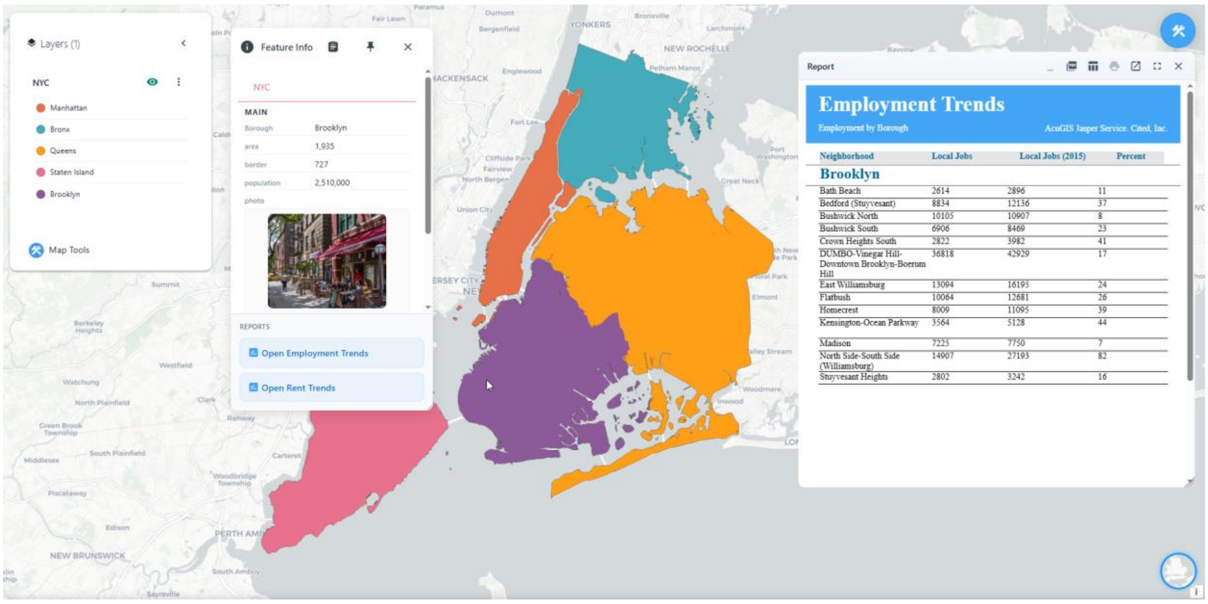

Published maps can integrate operational reporting directly within the user workflow.

Figure 11 — Jasper Reporting

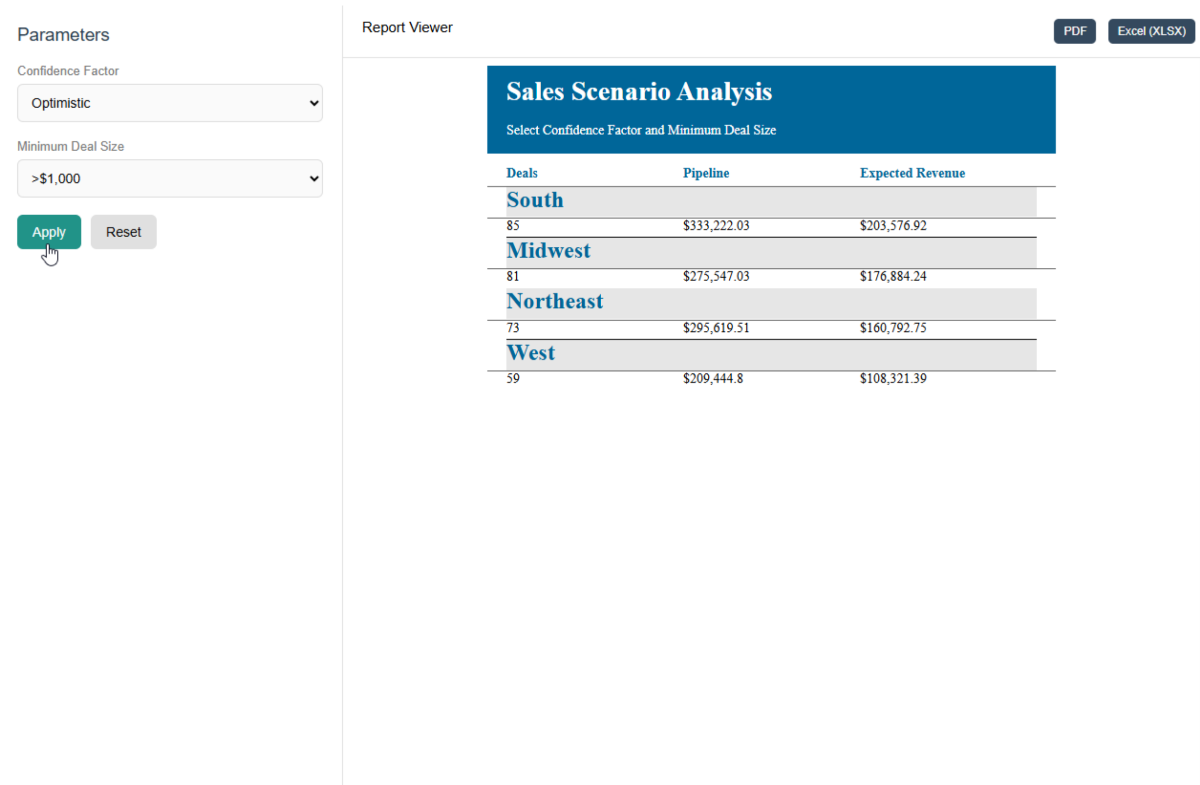

Jasper-based reports support interactive parameters, scheduled execution, and multiple export formats.

Optional Figures

`geo-etl-python.png`

Include only if you want to emphasize:

advanced users may perform Python preprocessing

Otherwise omit.

Reason:
The white paper should emphasize deterministic pipelines, not custom scripting.