geomatics

MDPI

Article

Integrating Geospatial Intelligence and Public Procurement Data for Fiscal Oversight: Design, Implementation, and Evaluation of a Geoportal for Public Infrastructure Monitoring

Andrés Felipe Cifuentes [1], Carlos Andrés Rodado [1], Eduardo José Tous [1], Álvaro José Giraldo [1], Rosse M. Villamil-Cañas [1], Jaime A. Restrepo-Carmona [2], Diana Yate [2] and Luis Fletscher [3,*]

[1] Contraloría General de la República, Carrera 69 No. 44-35, Bogota 111071, Colombia
[2] Rotorr—Motor de Innovación, Av. El Dorado No. 44a-40, Bogota 111071, Colombia
[3] Department of Electronic Engineering, Faculty of Engineering, Universidad de Antioquia, Cll. 67 No. 53-108, Medellin 050010, Colombia
* Correspondence: luis.fletscher@udea.edu.co

Abstract

The increasing availability of public-sector data creates opportunities for spatially enabled fiscal oversight, but procurement, project-monitoring, financial, and geographic information are commonly maintained in heterogeneous institutional systems. This work presents the design, implementation, deployment, and evaluation of a geospatial intelligence geoportal for public-infrastructure oversight in Colombia. Its contribution is architectural and integrative rather than algorithmic: source-specific ETL pipelines, explicit administrative-to-spatial relationships, thematic federation, standardized GIS publication, and multi-scale interaction are combined into a recurrent and traceable decision-support environment. The infrastructure pipeline processed approximately 258,528 AppUI source records into 25,897 analytical records. A polygon-intensive environmental module was used as a technical stress case for the shared geospatial infrastructure, where coordinated database, geometry, GIS-service, and client-side optimizations reduced mean feature-service loading time from 51.92 s to 1.48 s, while representative DISTINCT queries improved by 96–99%. The platform supports territorial, attribute-based, and project-level examination of spatial, contractual, temporal, financial, and execution-related information while preserving source traceability. An operational load-and-stress session analyzed 1054 HTTP requests with a 99.9% successful-request rate; however, the protocol was not a standardized high-concurrency benchmark and does not establish maximum concurrent-user capacity. The results provide evidence of technical feasibility and operational behavior in the evaluated public-sector deployment but do not establish improvements in audit outcomes, transparency, or decision quality. The architecture documents a potentially reusable implementation pattern for other public-sector monitoring contexts; its transferability and scalability require independent replication and standardized evaluation.

Keywords: geoportal; geospatial intelligence; Geographic Information Systems; spatial data integration; fiscal oversight; public infrastructure; public procurement; open government data; spatial decision support

Check for updates

Academic Editors: Enrico Corrado Borgogno Mondino and Caterina Balletti

Received: 12 August 2026
Revised: 16 October 2026
Accepted: 21 October 2026
Published: 1 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The increasing availability of public-sector data and advances in geospatial technologies are transforming the way governmental institutions collect, integrate, analyze, and

communicate information. Geoportals and Spatial Data Infrastructures (SDIs) have played an important role in this transformation by facilitating access to geographically referenced information distributed across heterogeneous organizations and information systems [1,2]. In parallel, Open Government Data (OGD) initiatives have expanded access to administrative information and created new opportunities for transparency, accountability, and data-driven governance [3–6].

However, increased data availability does not necessarily imply analytical interoperability. Public-sector information is frequently distributed across independent platforms, structured according to different data models, updated at different frequencies, and only partially associated with geographic information [3,4,6,7]. These characteristics constrain joint interpretation when the same public intervention must be examined simultaneously from contractual, financial, temporal, institutional, and territorial perspectives.

Geographic information provides a common analytical dimension across these heterogeneous records. From a geomatics perspective, the problem therefore extends beyond displaying infrastructure projects on a map: administrative records must be transformed, spatially enabled, validated, and published through a consistent geospatial environment while preserving traceability to their institutional sources.

Previous research has demonstrated the value of geoportals and geospatial decision-support platforms in environmental management, territorial planning, natural-resource monitoring, and policy-oriented decision support [7–9]. Geomatics technologies have also been applied to construction monitoring and government auditing, including UAV-based progress assessment, GIS-supported monitoring of public works, and institutional guidance on the use of geospatial information in audits [10–12]. These studies confirm the usefulness of spatial information for governmental monitoring, but they generally address specific analytical tasks rather than the recurrent integration of procurement, project-monitoring, fiscal, temporal, and geographic information within a common operational environment.

This integration problem is particularly relevant in Colombia. SECOP I and SECOP II are the national electronic public-procurement information systems, while AppUI provides institutional information used for monitoring public-investment projects. The Contraloría General de la República (CGR), Colombia's national fiscal-oversight institution, also maintains internal analytical and audit information. These sources differ in schema, identifiers, geographic representation, update frequency, access mechanisms, and data-governance requirements.

The institutional Geoportal examined in this study supports several fiscal-oversight domains, although the primary application scope of the paper is public-infrastructure monitoring. The infrastructure component integrates AppUI information, public-procurement records, and strategically prioritized major public works through a common geospatial interaction model. An environmental module focused on páramo ecosystems is used as a cross-domain technical stress case because its polygon-based datasets impose greater geometry-processing and GIS-service demands on the same underlying architecture.

Accordingly, this study examines how an implemented geospatial intelligence architecture can transform heterogeneous public-procurement, project-monitoring, fiscal, and geographic information into a traceable and updateable spatial decision-support environment. The study is organized around the following research questions:

RQ1. *How can heterogeneous procurement, project-monitoring, fiscal, and geospatial records be transformed into a traceable and updatable spatial analytical environment while preserving source-specific semantics?*

RQ2. *What technical performance effects result from combining database indexing, geometry optimization, GIS-service configuration, and client-side query optimization within the shared geospatial architecture?*

RQ3. *Can the implemented environment execute a representative fiscal-oversight workflow that jointly examines geographic, contractual, financial, temporal, and execution-related information from territorial to project level?*

The main contributions of this work are: (i) an operational integration pattern that recurrently transforms heterogeneous administrative and geographic information into spatially enabled analytical structures while preserving source traceability; (ii) a source-specific spatial-enablement and harmonization strategy based on explicit administrative-to-geographic relationships, validation, geometry processing, and multi-scale exploration; (iii) a thematic-federation architecture that preserves source-specific semantics while exposing analytical entities through standardized GIS services and a common interaction model; and (iv) an application-oriented evaluation based on production deployment, geospatial validation, geometry-complexity characterization, response-time measurements, load-and-stress testing, and an illustrative fiscal-oversight workflow.

The contribution is therefore architectural and integrative rather than algorithmic. The study does not propose a new geospatial algorithm or an automated mechanism for determining fiscal irregularities. Instead, it investigates how established geospatial and data-engineering technologies can be systematically combined into a recurrent operational architecture for fiscal-oversight analysis.

The remainder of the paper is organized as follows. Section 2 reviews related work and positions the research gap. Section 3 presents the implemented architecture, data-integration and spatial-enablement procedures, analytical model, and evaluation protocol. Section 4 reports deployment, geospatial validation, technical performance, load-and-stress results, analytical capabilities, and an illustrative fiscal-oversight workflow. Section 5 discusses the findings, contributions, and limitations in relation to the research questions. Section 6 summarizes the conclusions and future research priorities.

2. Related Work

Geoportals and Spatial Data Infrastructures (SDIs) are established mechanisms for discovering, accessing, integrating, and disseminating geographically referenced information across heterogeneous organizations and technical environments [1,2,13,14]. Their role has progressively evolved from spatial-data publication toward analytical and decision-support environments in which geospatial information is combined with administrative, environmental, socioeconomic, or operational data [9]. This evolution is particularly relevant in public-sector settings, where technological interoperability must be complemented by data governance, institutional coordination, and sustainable information-sharing arrangements [15].

2.1. Open Government Data and Transparency Platforms

Open Government Data (OGD) initiatives have expanded access to public-sector information and created opportunities for transparency, accountability, participation, and data-driven innovation [3–6]. However, public availability does not guarantee analytical interoperability. Differences in schemas, identifiers, update cycles, access mechanisms, and data quality can prevent records associated with the same public intervention from being examined jointly.

This limitation is especially relevant when procurement, project-monitoring, financial, and geographic information are distributed across independent systems. Associating administrative records with geographic entities provides a common analytical dimension through which these heterogeneous sources can be explored territorially. From this perspective, the challenge extends beyond data availability toward spatially enabled analytical reuse.

2.2. Geospatial Platforms for Governance and Decision Support

Web-based GIS and SDI developments have progressively combined spatial-data access, interactive visualization, analytical services, and decision-support capabilities. MapX, for example, integrates authoritative environmental information with visualization and analytical functionality [8], while Ran and Nedovic-Budic [9] distinguish conventional data-oriented infrastructures from online environments designed to support policy-oriented decision processes.

These approaches demonstrate the evolution of geomatics from information provision toward integrated decision support. Fiscal oversight introduces a different integration requirement: geographic information must be related to procurement, project-monitoring, financial, temporal, and audit-oriented records while preserving traceability to their original institutional systems.

Geospatial technologies have also been applied to infrastructure monitoring and public-sector auditing. GIS-supported public-works monitoring has demonstrated the usefulness of combining administrative and spatial information in audit activities [12], while INTOSAI has formally recognized the role of geospatial information in auditing through ISSAI 5540 [13]. Complementary approaches use UAV-based observation to assess construction progress [11], and megaproject research emphasizes the complexity and materiality of major public works [16].

These contributions demonstrate the value of spatial evidence for monitoring physical assets and audit activities. Their primary focus, however, is generally a specific monitoring or audit task rather than the recurrent integration of procurement, monitoring, financial, institutional, and geographic information within a common operational environment.

2.3. Integration of Geospatial and Contractual Data

Integrating procurement and geospatial information presents additional challenges because the underlying data models represent different types of entities. Procurement records are primarily transactional and administrative, whereas geospatial datasets are organized around geometries, coordinate reference systems, and spatial relationships. Connecting these domains therefore requires identifier matching, schema harmonization, spatial enablement, validation, metadata management, and interoperable service access.

OGC standards provide an important foundation for exchanging and consuming geographic information across heterogeneous systems [14], but interoperability standards alone do not determine how administrative records should be transformed into spatially queryable analytical entities. In the present study, contractual and execution records become objects of spatial analysis through explicit relationships with infrastructure projects or geographic entities while retaining traceability to the originating institutional sources.

2.4. Research Gap and Positioning of the Proposed Approach

The reviewed literature shows that the individual components required for geospatially enabled fiscal oversight are well established: geoportals and SDIs support spatial-data access and interoperability; OGD initiatives expose administrative information; decision-support infrastructures connect geospatial data with analytical processes; and audit-oriented applications demonstrate the usefulness of spatial evidence.

The research gap addressed here is therefore not the absence of GIS in public auditing. Rather, it concerns the limited documentation and evaluation of operational architectures that recurrently transform heterogeneous procurement, project-monitoring, fiscal, temporal, and geographic records into a traceable spatial analytical environment for fiscal-oversight workflows.

Existing approaches address important subsets of this problem. Table 1 summarizes their primary focus and clarifies the positioning of the present study.

Table 1. Comparison of representative related approaches and their relationship to the present study.

Reference	Application Domain	Primary Data Focus	Role of Geospatial Information	Limitation Relative to Present Study
[1]	Geoportals/SDI	Distributed spatial resources	Discovery and access	No fiscal-oversight or procurement integration
[8]	Natural resources	Authoritative environmental data	Integrated visualization and analysis	Different domain; no procurement/fiscal workflow
[9]	Planning/flood risk	Spatial planning and risk information	Policy-oriented decision support	No fiscal-oversight integration
[10]	National SDI governance	Institutional geospatial infrastructure	Governance and performance assessment	Not a project-level analytical architecture
[11]	Construction monitoring	Physical-progress observations	Acquisition and assessment of spatial evidence	No procurement/fiscal data integration
[12]	Public-works auditing	Audit and public-works information	GIS-supported monitoring	Task-oriented application; limited end-to-end integration

The distinctive contribution of the present work is therefore not a new GIS algorithm, but the end-to-end operational combination of recurrent administrative-data transformation, explicit spatial enablement, thematic federation, standardized GIS publication, and fiscal-oversight interaction in a production public-sector environment. The system does not automate the determination of fiscal irregularities, and the evaluation does not establish causal improvements in audit outcomes. Instead, the study evaluates whether heterogeneous institutional records can be transformed into an operational spatial decision-support environment and characterizes the technical behavior of that environment under representative use conditions.

3. Methodology

The research follows an implementation-oriented systems methodology in which architectural design, source-specific data transformation, spatial enablement, GIS publication, and operational evaluation are examined as an integrated workflow.

As was stated previously, the methodological focus is not on proposing a new spatial algorithm. Instead, it is on the reproducible transformation of administrative information into GIS-enabled analytical entities, the preservation of source traceability, the organization of heterogeneous domains through thematic federation, and the evaluation of the resulting architecture in an operational public-sector setting.

The implementation combines scheduled ETL processes, a spatially enabled SQL Server environment, ArcGIS Enterprise 11.4 services, web-based geo-viewers, role-based access, and separate presentation, application, geospatial-service, and data-processing components. The following subsections describe the design decisions and processing stages required to move from institutional source records to user-facing fiscal-oversight workflows.

3.1. System Architecture

The principal architectural decision was to separate source ingestion and transformation, analytical persistence, GIS publication, and user interaction. This separation was selected to reduce coupling with heterogeneous institutional systems and to allow thematic applications to evolve independently. The trade-off is increased orchestration complexity across ETL, database, GIS-service, and frontend components; however, this organization

preserves source ownership and avoids embedding source-specific reconciliation logic in the user-facing applications. Figure 1 summarizes this architecture, while Table 2 identifies the technologies used to implement each layer.

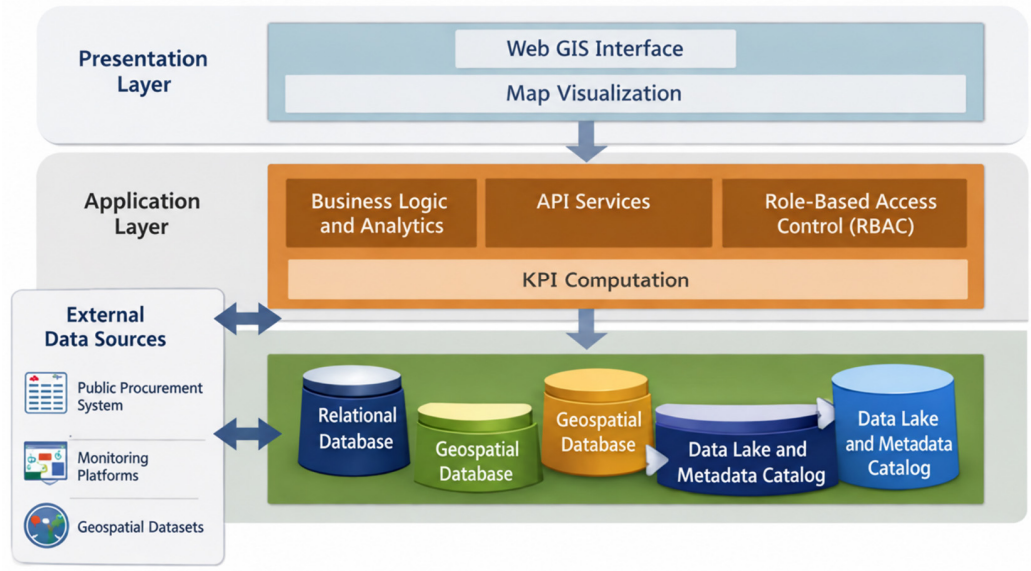


Figure 1. Layered architecture of the implemented geoportal, showing presentation, application, data, GIS-service, and integration components.

Table 2. Technological Components and Functions by System Layer.

Layer	Technology	Function
Frontend (Presentation)	Angular 19 (TypeScript)	Visualization of projects and indicators through maps and dashboards.
Application Layer (Backend)	Python 3.12/Django 5.0.14	Business logic, API orchestration, and authentication.
Data Layer	Microsoft SQL Server 2025 (Spatial)	Central repository for georeferenced, contractual, and financial data.
GIS Layer	ArcGIS Enterprise Cloud	Spatial rendering, geoprocessing, and thematic layer publication.
Integration Layer	Azure Data Factory/Databricks	ETL processes for extraction, transformation, and data harmonization.

Overall, the architecture combines modular application delivery, controlled data integration, spatially enabled persistence, and standardized GIS services. This organization is consistent with established SDI principles of distributed access and interoperability [1,2,7,14], while extending them toward an operational fiscal-oversight decision-support workflow.

3.2. Data Sources and Integration

The geoportal integrates heterogeneous information from public and institutional sources through scheduled, source-specific ingestion processes. The principal sources include SECOP I and II for public-procurement and contract information; AppUI for public-investment project monitoring; CGR internal repositories for fiscal, audit, and analytical information; and external public sources providing environmental, socioeconomic, and territorial reference data.

Data ingestion is implemented through API access, direct database connections, ETL workflows, and other source-specific interfaces according to the characteristics and update

frequency of each system. Azure Data Factory and Databricks support extraction, normalization, transformation, and loading into the analytical and geospatial environment. Rather than assuming real-time synchronization, the architecture applies scheduled recurrent ingestion according to source-specific update cycles, as summarized in Table 3.

Table 3. Main institutional and external data sources integrated in the CGR geoportal.

Source	Type	Frequency	Integration Method	Responsible Entity
SECOP I–II	Procurement/Contracts	Daily	API/ETL	Colombia Compra Eficiente
AppUI	Investment Progress/public works progress	Weekly	ETL/Databricks	DNP
CGR Internal Data	Audits/Indicators	Weekly	Direct DB	CGR
Environmental Datasets	Land cover/Deforestation	Monthly	FTP/REST	IDEAM
Open Data Portals	Socioeconomic/Territorial	Monthly	API	Gov.co

SECOP I–II and official environmental and open-data sources are publicly accessible, whereas AppUI and CGR internal repositories are subject to institutional access and data-governance restrictions.

Geographic publication follows institutional reference-system conventions. MAGNA-SIRGAS/Origen-Nacional (EPSG:9377) is used for institutional spatial data, while WGS84 (EPSG:4326) is supported for interoperable layers. Spatial Reference Identifiers (SRIDs) are verified before publication; this choice follows institutional publication conventions rather than a comparative CRS experiment.

3.3. Data Integration, Georeferencing, and Spatial Harmonization

The conversion of heterogeneous institutional records into spatially enabled analytical information is a central component of the implemented geoportal. Because the source systems differ in structure, granularity, geographic representation, and completeness, raw source tables are not directly exposed to the visualization layer. Instead, source-specific ETL processes perform profiling, cleaning, filtering, restructuring, and standardization before records are incorporated into the analytical and geospatial environment.

Transformation, spatial association, validation, GIS publication, and interactive visualization are treated as separate processing stages. This separation preserves traceability to the original institutional systems and prevents the frontend from performing ad hoc geocoding or source-record reconciliation during user interaction.

A central design choice was to use source-specific deterministic spatial relationships rather than a universal geocoding mechanism. This preserves the semantics and stable identifiers available in each institutional domain and avoids probabilistic location inference. The trade-off is that spatial-enablement rules must be maintained separately for sources with different geographic representations, and records without adequate source-level spatial references cannot be retrospectively assigned a quantitative matching-confidence score.

3.3.1. Source-Specific Data Transformation

The geoportal applies different transformation rules according to the analytical purpose and granularity of each thematic domain. Public-infrastructure monitoring integrates AppUI project information, public-procurement records, and strategically prioritized major public works, while the Páramos module independently processes investment, contracting, and audit-finding records.

ETL therefore represents more than source replication. Filtering, restructuring, consolidation, and analytical-granularity rules transform source records into structures suitable for geospatial analysis. For AppUI, approximately 258,528 source records produce 25,897 analytical records. In the Páramos domain, 462 investment records produce 299 analytical records, 16,775 contracting records produce 7863 analytical records, and 77 audit-finding records generate 404 analytical records because the final representation uses a different analytical granularity. These counts should therefore be interpreted as the result of source-specific transformation rather than as simple retained-versus-discarded observations.

3.3.2. Spatial Enablement and Geometry Harmonization

After ETL processing, analytical records are incorporated into dedicated spatial structures rather than queried directly from the original institutional systems. Administrative information becomes spatially queryable through explicit relationships with validated geographic entities, while the original source attributes and identifiers are preserved for traceability.

The Páramos module provides a representative example. Investment, contracting, and audit records are associated with the corresponding geographic unit through a deterministic relational-spatial join using the stable institutional identifiers paramo_id and paramo_objectid. Spatial enablement is therefore source-specific and identifier-based rather than dependent on text-derived or probabilistic geocoding.

Geometry harmonization is incorporated into the recurrent loading process. Polygon geometries are simplified using Shape.Reduce(500) during data preparation, and geographic extents are recalculated during subsequent ETL executions. The original production layers contained geometries with up to 2530 vertices for contracting and investment records and up to 4934 vertices for audit-finding records. A second, scale-dependent simplification mechanism is applied during visualization, using lower geometric detail at broader map scales and full detail for local inspection.

Consequently, spatial harmonization is a recurrent component of the ETL workflow rather than a one-time manual operation. The available technical evidence characterizes geometry complexity and its performance effects, but does not provide an independent geometric-fidelity benchmark based on metrics such as area error, Hausdorff distance, or centroid displacement; this limitation is addressed in Section 5.4.

3.3.3. Query Optimization and GIS Publication

Spatial structures are stored in the enterprise geodatabase and published through ArcGIS Enterprise as standardized GIS services. This separates institutional source models from the analytical structures consumed by the geoportal and allows source-specific transformation rules to evolve without requiring changes in the visualization layer.

Database structures are optimized according to the query patterns generated by the geo-viewers. Nonclustered indexes are applied to frequently used spatial-linkage, temporal, and thematic-filter attributes to improve joins and interactive queries. GIS-service and client-side mechanisms complement these optimizations by limiting transferred attributes, loading thematic layers on demand, and applying scale-dependent geometry generalization.

The implemented geoportal exposes four principal thematic viewers—AppUI, Megaobras, Public Procurement, and Páramos—through a common navigation and service-publication model. Although each domain retains its own semantics and data structures, the processing sequence is consistent:

source records → source-specific transformation → spatial association and validation → analytical geospatial structures → GIS-service publication → interactive geo-viewer.

This sequence preserves separation between institutional sources, analytical processing, GIS publication, and frontend interaction while providing a common geospatial access model across heterogeneous fiscal-oversight domains.

3.4. Analytical Data Model and Query Processing

The analytical architecture supports joint examination of fiscal, contractual, temporal, and geographic information without imposing a single universal physical schema. The implementation does not rely on a cross-domain ontology. Instead, it follows a thematic-federation approach in which source-specific schemas and transformation logic are preserved, while interoperability is achieved through explicit spatial relationships, standardized validation, common GIS-publication mechanisms, and shared interaction patterns.

Each spatially enabled analytical record may combine institutional and contractual attributes, temporal information, fiscal or execution-related indicators, and the geographic entity through which the record is represented. Independent thematic datasets are maintained according to the semantics and granularity of each domain. This preserves source-specific meaning and avoids forcing heterogeneous records into a large denormalized model. The corresponding trade-off is that cross-domain interoperability depends on explicit transformation and spatial-linkage rules rather than on a universal ontology.

3.4.1. Thematic Federation and Spatial Relationships

The geoportal maintains separate analytical structures for the principal fiscal-oversight domains. For infrastructure monitoring, AppUI represents institutionally reported investment projects, Megaobras contains strategically prioritized public works, and the Public Procurement model provides the contractual dimension. The Páramos domain maintains separate investment, contracting, and audit-finding structures because these datasets correspond to different units of observation and analytical granularities.

Across these thematic structures, the common analytical principle is the explicit relationship between the administrative entity represented by the source record and the spatial entity used for geographic analysis. Depending on the domain, the spatial entity may correspond to an infrastructure project, administrative territory, or environmental geographic unit. Once this association is established, the corresponding geographic feature becomes the entry point for retrieving contractual, financial, temporal, institutional, execution-related, or audit attributes.

Thus, commonality across domains is provided by spatial enablement, standardized GIS publication, filtering mechanisms, and shared interaction patterns rather than by forcing semantically different records into a single schema.

3.4.2. Query and Service Design

The physical analytical structures are optimized according to the query patterns generated by the thematic geo-viewers. Frequently queried spatial-linkage, temporal, and thematic attributes are indexed at the database level rather than relying exclusively on application-side filtering. This design reduces query and service latency, although indexes and dependent views must be reviewed when source structures or interaction patterns change.

Database-level optimization is complemented by service and client-side mechanisms that restrict transferred attributes, adapt query sizes to the analytical component, and avoid unnecessary loading of thematic layers. Together, these mechanisms reduce the amount of information processed and transferred during interactive geographic analysis.

Spatially enabled analytical structures are not exposed directly to client applications. After processing and validation, they are published through ArcGIS Enterprise feature

or map services. This service-oriented organization separates responsibilities among the analytical database, GIS-service layer, and geo-viewers and helps preserve consistency between processed data and published spatial resources.

Traceability to institutional sources is maintained throughout the processing chain because the geoportal operates as a decision-support layer rather than as a replacement for the originating information systems. Source records remain under the responsibility of their institutional systems, while recurrent ETL processes transform updated information into the corresponding analytical structures.

This separation allows source-specific transformation and spatial-processing rules to evolve without requiring equivalent changes in the frontend interaction logic. The resulting architecture therefore combines source preservation, thematic analytical modeling, spatial enablement, optimized query access, and standardized GIS publication while retaining domain-specific semantics.

3.5. Geospatial Services, User Interaction, and Operational Architecture

Spatially enabled analytical datasets are exposed through ArcGIS Enterprise services and consumed by thematic web-based geo-viewers, avoiding direct frontend access to the original institutional databases. The implemented environment comprises four principal viewers—AppUI, Megaobras, Public Procurement, and Páramos—integrated through a common navigation model while retaining independent thematic data structures.

From the user perspective, the shared interaction model combines multi-scale geographic navigation, thematic-layer selection, spatial and attribute-based filtering, project-level inspection, and visualization of fiscal or execution-related indicators. This allows users to move from territorial patterns to the associated contractual, temporal, financial, execution, or audit information while preserving the semantics of each thematic domain.

Access to information and functionality is controlled according to institutional requirements. Public consultation is separated from authenticated institutional access, and role-based authorization governs functions involving restricted datasets or administrative privileges. Security, logging, data protection, and governance are therefore treated as cross-cutting operational controls rather than as independent application modules.

The production deployment separates user-facing applications from backend and GIS-service components and supports modular evolution of the thematic geo-viewers. Performance is addressed through coordinated database, geometry, GIS-service, and client-side optimizations, while detailed performance and load-and-stress evaluation procedures are reported in Section 3.6 and their results in Section 4. Infrastructure-specific addresses and security-sensitive configuration details are intentionally omitted.

3.6. Evaluation Protocol

The implemented geoportal was evaluated through five complementary dimensions: (i) functional validation, (ii) geospatial data-quality and geometry validation, (iii) comparative technical-performance evaluation, (iv) operational load-and-stress evaluation, and (v) an illustrative fiscal-oversight workflow. The protocol focuses on the behavior of the deployed system under conditions representative of institutional operation rather than on a conceptual prototype.

3.6.1. Functional Validation

Functional validation examined the end-to-end operation of the deployed geo-viewers, including geographic navigation, thematic-layer visualization, spatial and attribute-based filtering, project-level queries, retrieval of contractual and execution-related information, and differentiated access to analytical components. Validation covered the complete interac-

tion sequence between the user interface, GIS services, analytical datasets, and underlying data structures; successful completion of the expected workflow and consistent retrieval of spatial and attribute information constituted the validation criterion.

3.6.2. Geospatial Data-Quality and Geometry Validation

Spatial datasets were evaluated before GIS publication using record-count verification, geometry-validity checks, detection of null geometries, SRID consistency, duplicate-key detection, attribute-type and range checks, and correspondence with institutional reference layers. For the Páramos stress case, geometry complexity was additionally characterized through vertex counts, and deterministic administrative-to-spatial linkage was verified through the institutional identifiers used to associate fiscal records with geographic units.

The available engineering records do not retain a complete historical record of unmatched eligible source records, nor was an independent positional ground-truth dataset preserved during the original operational deployment. Consequently, geocoding success, spatial-linkage completeness, and independent positional-accuracy rates cannot be retrospectively reconstructed. This absence is a limitation of the retained evaluation records rather than a methodological feature or contribution of the proposed integration approach. The present evaluation therefore reports only those spatial-quality controls that can be supported by the available evidence: geometry validity, null-geometry detection, SRID consistency, duplicate-key controls, deterministic identifier-based linkage, and geometry-complexity characterization.

3.6.3. Technical Performance Evaluation Protocol

Technical performance was evaluated using the Páramos geo-viewer as a demanding operational case for the shared geospatial infrastructure because its polygon-based layers impose greater geometry-processing and GIS-service demands than the predominantly project-oriented infrastructure datasets.

HTTP Archive (HAR) network captures were collected at three stages of the optimization process: (i) baseline configuration, (ii) post-indexing configuration, and (iii) the final optimized state after coordinated database, geometry, GIS-service, and application-level interventions. The evaluated indicators included feature-service loading time, execution time of representative thematic-filter queries, investment-data query time, and observed HTTP request failures.

Relative performance improvement was calculated as

$$\text{Improvement}(\%) = \frac{T_{baseline} - T_{optimized}}{T_{baseline}} * 100,$$

where $T_{baseline}$ is the observed response time before optimization and $T_{optimized}$ is the corresponding response time after the complete optimization process.

Because these measurements were obtained during iterative operational optimization, they are interpreted as comparative performance evidence rather than as a standardized platform-capacity benchmark.

3.6.4. Operational Load-And-Stress Evaluation

An instrumented operational load-and-stress session was conducted on the deployed geoportal infrastructure to complement the comparative HAR-based performance evaluation. The session generated representative application demand involving the Public Procurement and Megaobras modules and was used to observe HTTP behavior and infrastructure-resource utilization.

Infrastructure monitoring included CPU utilization and available memory, while HTTP traffic analysis considered request status codes, mean and maximum response times, and

transferred data volume. These measurements were used to identify potential resource bottlenecks, request failures, and module-specific response-time anomalies.

The session provides quantitative evidence of system behavior under the specific evaluated demand. It was not designed as a standardized high-concurrency benchmark: the protocol did not prescribe virtual-user levels, a progressive concurrency ramp, repeated load stages, throughput targets, or a maximum-capacity threshold. Its results are therefore interpreted as operational stress evidence and not as a measurement of maximum concurrent-user capacity or horizontal scalability.

3.6.5. Fiscal-Oversight Workflow Evaluation Protocol

Finally, practical applicability was examined through an illustrative fiscal-oversight workflow in which an institutional user selects a geographic area, filters infrastructure projects according to contractual and execution-related attributes, inspects project-level information, and jointly examines spatial location together with physical and financial indicators.

The purpose of the workflow is to verify whether the integrated environment supports the information-retrieval and geographic-contextualization sequence required for subsequent audit examination. It is not a controlled test of audit efficiency, anomaly-detection accuracy, or fiscal-finding validity.

4. Results

The geoportal was deployed as an operational geospatial environment integrating public-procurement, institutional project-monitoring, fiscal, and geographic information. Results are presented in terms of (i) deployment and data integration, (ii) geospatial validation and geometry complexity, (iii) technical performance and load-and-stress behavior, (iv) geospatial analytical capabilities, and (v) an illustrative fiscal-oversight workflow. Public infrastructure constitutes the primary application domain, while the Páramos module is used as a cross-domain technical stress case because its polygon-based datasets impose greater geometry-processing and GIS-service demands on the shared infrastructure.

4.1. Deployment and Integrated Infrastructure Data Environment

The geoportal was implemented and deployed as a spatially enabled decision-support environment for fiscal oversight of public infrastructure in Colombia. The infrastructure component integrates AppUI, public-procurement records, strategically prioritized major public works, and geographic reference information through the recurrent ETL, validation, and GIS-publication workflow described in Section 3.

Source datasets differed substantially in volume and analytical granularity. Approximately 258,528 AppUI source records were transformed into 25,897 analytical records, whereas the Megaobras dataset retained its 107 strategically prioritized projects. These differences reflect source-specific filtering, consolidation, restructuring, and standardization rather than simple record elimination.

The resulting analytical structures preserve source-specific semantics while exposing institutional, contractual, temporal, execution-related, and geographic information through a common interaction model. Users can therefore explore infrastructure information from national and territorial views to individual project records without directly querying or reconciling the original institutional systems.

4.2. Geospatial Validation and Geometry-Complexity Results

The recurrent publication workflow applied record-count verification, geometry-validity and empty-geometry checks, SRID consistency, duplicate-key detection, attribute-coherence controls, and comparison with institutional reference layers before GIS publi-

cation. The available engineering records confirm that these controls were applied, but they do not preserve a complete per-record history of rejected, unmatched, or spatially corrected records. Consequently, quantitative geocoding-success, linkage-completeness, and independent positional-accuracy rates cannot be retrospectively reconstructed. This limitation reflects the historical data records available for evaluation and should not be interpreted as a methodological contribution of the system.

The Páramos module provides quantitative evidence of dataset scale and geometry complexity. After spatial migration, the production geodatabase contained 325 investment records, 6901 contracting records, and 404 audit-finding records. These values correspond to the published spatial stage and therefore differ from the upstream transformation counts reported in Section 3.3.

Before geometry optimization, investment and contracting polygons reached up to 2530 vertices per geometry, while audit-finding geometries reached 4934 vertices. Average complexity was 128, 135, and 461 vertices per geometry, respectively (Table 4). These characteristics support the use of the Páramos module as the geometry-intensive case for evaluating the shared GIS infrastructure.

Table 4. Geometry-complexity characteristics of the Páramos production layers before service-level optimization.

Spatial Analytical Layer	Published Records	Average Vertices/Geometry	Maximum Vertices/Geometry	Spatial Association
Investment	325	128	2530	paramo_id/paramo_objectid
Contracting	6901	135	2530	paramo_id/paramo_objectid
Audit findings	404	461	4934	paramo_id/paramo_objectid

4.3. Technical Performance Evaluation of the Shared Geospatial Architecture

Technical performance was evaluated using the Páramos geo-viewer as a demanding operational scenario for the shared ArcGIS-based infrastructure. This module contains polygon geometries with substantially greater geometric complexity than the point- and project-oriented infrastructure datasets and was therefore used to assess the effects of database indexing, geometry generalization, GIS-service configuration, caching, and front-end query optimization.

4.3.1. Comparative Optimization Results

Comparative performance measurements from the Páramos module showed substantial reductions after coordinated database, geometry, GIS-service, and application-level optimization. Table 5 reports the baseline, post-indexing, and final optimized measurements obtained from HTTP Archive captures.

The mean loading time of the principal Páramos feature service decreased from 51.917 s to 1.482 s, representing a 97% reduction. Representative DISTINCT queries showed similarly large improvements: the investment query decreased from 34.087 s to 0.292 s and the audit-findings query from 21.657 s to 0.885 s.

The intermediate measurements also indicate that database indexing alone did not explain the complete improvement. Feature-service loading decreased from 51.917 s to 8.003 s after indexing and subsequently to 1.482 s after the complete optimization process, indicating a multi-layer performance bottleneck rather than an exclusively database-level limitation.

One investment-data query showed a less monotonic pattern, decreasing from 13.418 s at baseline to 0.240 s after indexing and reaching 2.715 s in the final measurement. This value was retained rather than excluded because it illustrates the variability of operational

measurements and the influence of factors such as cache state and service behavior. Overall, the results should therefore be interpreted as comparative operational evidence rather than as a standardized cross-platform benchmark.

Table 5. Response-time measurements before and after optimization.

Performance Indicator	Baseline (ms)	Post-Indexing (ms)	Optimized State (ms)	Reduction
Páramos feature-service load—mean	51.917	8.003	1.482	97%
Páramos feature-service load—maximum	63.650	29.383	4.768	93%
Audit-findings DISTINCT query—maximum	21.657	17.376	885	96%
Investment DISTINCT query—maximum	34.087	292	292	99%
Investment query—maximum	13.418	240	2.715 *	80%

* The optimized-state value corresponds to a cold-cache execution; an approximately 240 ms response was observed under warm-cache conditions.

4.3.2. Operational Load-And-Stress Results

The operational load-and-stress evaluation analyzed 1054 HTTP requests generated during representative interaction with the deployed geoportal. The global mean response time was 181.78 ms, with a 99.9% successful-request rate. Of the analyzed requests, 1051 returned HTTP 200, one returned HTTP 204, and two were cancelled or affected by network-level errors; no HTTP 404 or 500 responses were observed.

Performance differed between the evaluated application modules. Public Procurement showed a mean response time of 70 ms and a maximum of 983 ms, whereas Megaobras showed a mean of 281 ms and a maximum of 16.081 s. The latter isolated peak indicates that module-specific latency may occur even when aggregate response behavior remains stable and infrastructure resources are not saturated.

Infrastructure monitoring showed no exhaustion of CPU or memory resources during the evaluated session. The most heavily loaded service component reached a CPU peak of approximately 68%, while the database/GIS component remained substantially below saturation and the frontend showed low utilization. These results (Table 6) provide evidence of stable behavior under the evaluated load conditions but do not establish maximum throughput or high-concurrency capacity because the test documentation does not define standardized virtual-user levels or a progressive concurrency ramp.

Table 6. Summary of the operational load-and-stress evaluation.

Metric	Result
Analyzed HTTP requests	1054
Global mean response time	181.78 ms
Successful-request rate	99.9%
HTTP 200	1051
HTTP 204	1
Cancelled/network-level errors	2
HTTP 404/500	0
Public Procurement mean/max	70/983 ms
Megaobras mean/max	281/16.081 ms

4.4. Geospatial Visualization and Analytical Capabilities

The implemented geoportal supports multi-scale exploration of public-infrastructure information from national and territorial views to individual project records. Map-based navigation is combined with geographic and attribute filtering according to variables such as sector, execution status, temporal criteria, and contractual characteristics.

The interaction model supports three complementary analytical levels: territorial exploration of project distributions, attribute-based filtering of contractual and execution information, and project-level examination of spatial location together with associated administrative indicators. Figure 2 illustrates these capabilities within the implemented interface.

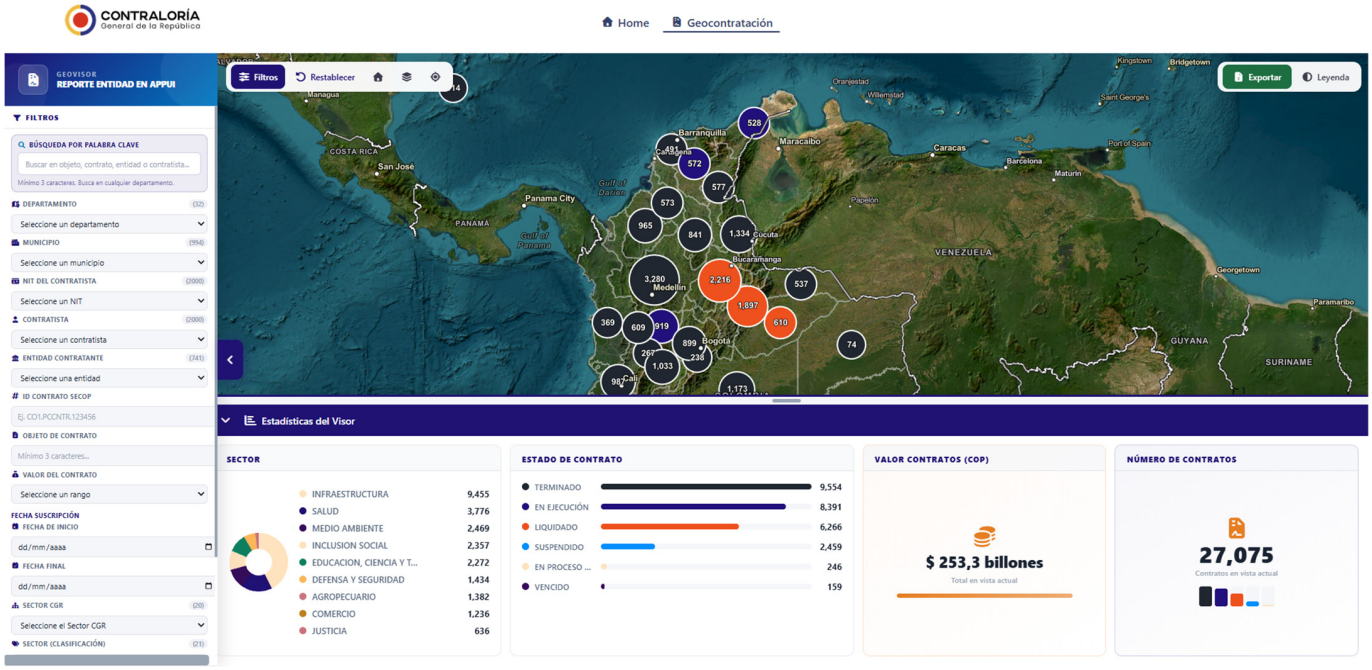


Figure 2. Implemented geoportal interface illustrating integrated cartographic navigation, thematic filters, project-level attributes, and analytical indicators.

4.5. Illustrative Fiscal-Oversight Workflow

An illustrative fiscal-oversight workflow was examined to determine whether the integrated environment supports an end-to-end project-monitoring sequence. The scenario involved an institutional user selecting an administrative territory, filtering infrastructure projects by sector and execution status, geographically examining the resulting records, and jointly inspecting available physical and financial execution indicators.

The workflow illustrates the operational integration of geographic, contractual, temporal, and execution-related information within a single interaction sequence. Differences between reported physical and financial progress can be geographically contextualized and inspected as inputs for subsequent professional examination. However, the workflow was not designed to measure whether the geoportal improves the accuracy, speed, or effectiveness of auditing activities.

Because the available evaluation records do not preserve a fixed experimental sample, task-completion times, comparative auditor performance, or adjudicated audit outcomes, this workflow provides evidence of workflow operability and information integration only. It does not provide quantitative evidence of audit effectiveness, anomaly-detection performance, or fiscal-finding validity. Determination of fiscal irregularities remains part of the institutional audit process.

5. Discussion

The results demonstrate the technical feasibility of transforming heterogeneous procurement, project-monitoring, fiscal, temporal, and geographic information into a traceable geospatial environment for public-infrastructure oversight. The contribution is primarily architectural and operational: established data-engineering and geospatial technologies are combined through source-specific transformation, explicit spatial enablement, thematic federation, standardized GIS publication, and coordinated performance optimization.

The evidence supports three conclusions aligned with the research questions. First, heterogeneous institutional information can be recurrently spatially enabled while preserving source-specific semantics and traceability (RQ1). Second, coordinated optimization across database, geometry, GIS-service, and application layers can substantially improve operational response times and support stable behavior under the evaluated load conditions (RQ2). Third, the deployed environment can support an end-to-end fiscal-oversight information workflow from territorial exploration to project-level inspection (RQ3). These findings concern technical integration and workflow support and should not be interpreted as evidence of improved audit outcomes, transparency, or decision quality.

5.1. Geomatics Contribution and Comparison with Existing Approaches

Regarding RQ1, the implemented approach demonstrates that heterogeneous administrative and geospatial information can be integrated without imposing a universal ontology or a single denormalized schema. Source-specific transformation and semantics are preserved, while interoperability is achieved through explicit administrative-to-spatial relationships, standardized validation, recurrent publication, and a shared geographic interaction model.

This approach extends established geoportal and SDI principles of spatial-data access and interoperability [1,2,7,15] toward an operational fiscal-oversight decision-support setting. Platforms such as MapX and online decision-support infrastructures similarly illustrate the evolution from data publication toward integrated analytical environments [8,9], but the present implementation applies these principles to the recurrent integration of procurement, project-monitoring, fiscal, and geographic information.

The use of thematic federation represents an explicit design trade-off. It preserves source semantics and allows thematic datasets to evolve independently, but requires maintained transformation, linkage, validation, and update rules. This contrasts with approaches based on a universal ontology and constitutes one of the principal architectural choices of the implemented system.

Regarding RQ2, the performance results indicate that no single intervention explains the observed improvements. Database indexing generated substantial query-time reductions, while final service performance also depended on geometry generalization, GIS-service configuration, caching, selective data retrieval, and client-side request optimization. The operational load-and-stress evaluation further showed stable aggregate behavior under the tested operational demand, although it does not establish the platform's maximum concurrency capacity.

5.2. Contribution to Open Government Data and Transparency

The implementation illustrates how a geospatial integration layer can increase the analytical usability of distributed public-sector information without replacing the originating institutional systems. Geographic location provides a common analytical reference through which procurement, project-monitoring, and other administrative information can be jointly explored while retaining source traceability.

This capability may support information accessibility and institutional analysis, but the present study does not measure changes in transparency, citizen participation, or decision quality. These should therefore be considered potential downstream effects rather than demonstrated outcomes and require dedicated user-centered and longitudinal evaluation.

5.3. Implications for Fiscal Oversight and Governance

Regarding RQ3, the illustrative workflow demonstrates that the deployed environment supports a continuous interaction sequence in which users can select a territory, filter projects, inspect individual records, and geographically contextualize contractual, physical, and financial information. This reduces the need to examine the same information through disconnected tabular and geographic environments.

The workflow nevertheless demonstrates information integration and contextualization, not improved audit effectiveness. No controlled comparison was performed to determine whether auditors complete tasks faster, identify more valid findings, or make better decisions using the geoportal. Likewise, the platform does not automatically classify discrepancies as fiscal irregularities or perform validated anomaly detection.

The system should therefore be interpreted as a decision-support environment that organizes and spatially contextualizes evidence for subsequent professional examinations. More advanced spatio-temporal, prioritization, or risk-oriented methods could be implemented on top of this architecture, but their validity would need to be assessed against adjudicated audit outcomes.

5.4. Limitations of the Study

Several limitations should be considered when interpreting the results.

First, analytical reliability depends on the quality, completeness, temporal consistency, and semantics of the originating datasets. Technical publication controls cannot correct errors or omissions already present in institutional source records.

Second, the available documentation does not contain a complete quantitative audit of spatial-linkage completeness or an independent positional ground-truth dataset. Consequently, the study reports deterministic linkage procedures, geometry-validity controls, SRID consistency, and geometry-complexity indicators, but not geocoding-success or positional-accuracy statistics. The geometric fidelity of Shape.Reduce(500) was also not evaluated using area error, Hausdorff distance, or centroid displacement. Future deployments should therefore incorporate explicit logging of matched, unmatched, corrected, and rejected spatial records so that these spatial-quality indicators can be calculated prospectively.

Third, the current analytical capabilities are predominantly descriptive and exploratory. The platform supports spatial visualization, filtering, geographic aggregation, and inspection of project indicators but does not implement validated predictive risk models or automated anomaly-detection algorithms.

Fourth, although an operational load-and-stress evaluation was performed, the available test protocol does not define standardized virtual-user levels, progressive concurrency ramps, repeated load scenarios, or a maximum-capacity threshold. The reported results therefore characterize operational behavior under the evaluated load conditions rather than providing a reproducible maximum-throughput or horizontal-scalability benchmark.

Finally, the study did not conduct a controlled user-centered comparison of audit efficiency, usability, transparency perception, or decision quality. Claims are therefore limited to spatial integration, technical feasibility, operational performance, and workflow support.

5.5. Future Research Directions

Future work should first strengthen geospatial-quality assessment by quantifying spatial-linkage completeness, positional consistency, semantic correspondence between administrative and geographic records, and the geometric fidelity of simplification operations.

A second priority is standardized scalability evaluation using predefined virtual-user levels, request mixes, concurrency ramps, repeated trials, latency distributions, throughput, and failure rates. These experiments would complement the controlled operational load-and-stress results reported here and support assessment of horizontal scaling and high-availability configurations.

Finally, the architecture provides a basis for both advanced analytical methods and user-centered evaluation. Future studies could examine spatio-temporal patterns, clustering, risk prioritization, or machine-learning approaches validated against adjudicated audit outcomes; evaluate retrieval time, usability, analytical effectiveness, and decision quality with auditors; and replicate the architecture in other institutions and geographic contexts to assess transferability beyond the Colombian fiscal-control environment.

6. Conclusions

This study presented the design, implementation, deployment, and evaluation of a geospatial intelligence geoportal for fiscal oversight of public-infrastructure projects in Colombia. Regarding RQ1, the implementation demonstrates that heterogeneous procurement, project-monitoring, fiscal, temporal, and geographic information can be recurrently transformed into spatially enabled analytical structures while preserving source-specific semantics and traceability.

From a geomatics perspective, the central contribution is the explicit spatial enablement of administrative information through validated geographic relationships, thematic federation, standardized GIS publication, and multi-scale exploration. The Páramos stress case further illustrates how geometry complexity affects operational behavior in a shared geospatial architecture and how persistent and scale-dependent simplification can be incorporated into recurrent processing.

Regarding RQ2, coordinated optimization across database, geometry, GIS-service, and client layers produced substantial response-time reductions: mean loading time for the principal Páramos feature service decreased from 51.917 s to 1.482 s, while representative DISTINCT queries improved by 96–99%. In addition, the operational load-and-stress evaluation recorded a 99.9% successful-request rate under the specific evaluated demand. Because the protocol was not a standardized high-concurrency benchmark, this result characterizes the observed session but does not establish maximum throughput, concurrent-user capacity, or horizontal scalability.

Regarding RQ3, the implemented environment supports an end-to-end workflow in which users can move from territorial selection and attribute filtering to project-level examination of geographic, contractual, financial, temporal, and execution-related information. The workflow provides integrated decision support but does not automate the determination of fiscal irregularities or establish improved audit effectiveness.

The main remaining limitations concern independent positional and spatial-linkage quality assessment, geometric-fidelity evaluation, standardized scalability testing, and formal user-centered studies. These areas define the next empirical steps required to evaluate the architecture beyond technical feasibility and operational deployment.

Overall, the implementation documents an architectural pattern that may be reusable for spatially enabled fiscal oversight: preserving source semantics, transforming records through recurrent ETL, establishing explicit administrative-to-spatial relationships, validating and publishing standardized GIS services, and exposing them through a common

analytical interaction model. Its transferability, scalability, and effects on audit practice require independent replication and controlled evaluation in other institutional settings.

Author Contributions: Conceptualization, A.F.C., C.A.R., D.Y., E.J.T., J.A.R.-C. and R.M.V.-C.; methodology, C.A.R., E.J.T., R.M.V.-C., L.F. and D.Y.; software, C.A.R., R.M.V.-C., E.J.T., D.Y. and L.F.; validation, C.A.R., R.M.V.-C. and E.J.T.; formal analysis, D.Y. and L.F.; investigation, A.F.C., C.A.R., E.J.T., J.A.R.-C., D.Y., L.F. and R.M.V.-C.; resources, A.F.C., J.A.R.-C. and R.M.V.-C.; data curation, C.A.R., E.J.T., D.Y. and R.M.V.-C.; writing—original draft preparation, D.Y. and L.F.; writing—review and editing, A.F.C., C.A.R., Á.J.G., J.A.R.-C., D.Y., L.F. and R.M.V.-C.; visualization, D.Y.; supervision, A.F.C., J.A.R.-C., L.F. and R.M.V.-C.; project administration, A.F.C., J.A.R.-C. and L.F.; funding acquisition, A.F.C., J.A.R.-C., L.F. and R.M.V.-C. All authors have read and agreed to the published version of the manuscript.

Funding: This work was funded by the Contraloría General de la República (CGR) under Contract No. CGR-407-2024, executed in collaboration with Universidad de Antioquia (UdeA).

Institutional Review Board Statement: Based on the nature of the data used in this paper, such approval was not required. We confirm that all data employed in this research fully comply with the guidelines established in the System for the Governance of Data and Information of the Contraloría General de la República, as defined in the Internal Resolution OGZ-0768-2020.

Informed Consent Statement: Waived as based on the nature of the data used in this paper; such approval was not required.

Data Availability Statement: The study combines public and institutionally restricted information. Public-procurement data originate from SECOP I and SECOP II, and public geographic reference datasets are obtained from official Colombian geospatial sources. AppUI and additional CGR project-monitoring, analytical, and audit-related information are subject to institutional access and data-governance restrictions and cannot be publicly redistributed by the authors. Processed datasets may be made available upon reasonable request to the corresponding author, subject to authorization by the data-owning institutions. To support methodological reproducibility without disclosing restricted infrastructure information, the manuscript documents the source-to-analytical workflow, spatial-linkage logic, validation criteria, geometry-processing strategy, and performance-evaluation protocol used in the study. No public repository of derived institutional artifacts is provided because their release remains subject to authorization by the data-owning institutions.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT (GPT-5.6, OpenAI, San Francisco, CA, USA) to assist with language editing, manuscript organization, and improvement of textual clarity. The authors reviewed and edited all generated content and take full responsibility for the content of the publication.

Conflicts of Interest: Authors Jaime A. Restrepo-Carmona and Diana Yate were employed by the company Rotorr Motor de Innovacion. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as potential conflicts of interest.

References

1. Maguire, D.J.; Longley, P.A. The emergence of geoportals and their role in spatial data infrastructures. *Comput. Environ. Urban Syst.* **2005**, *29*, 3–14. [CrossRef]
2. Rajabifard, A.; Feeney, M.-E.F.; Williamson, I.P. Future directions for SDI development. *Int. J. Appl. Earth Obs. Geoinf.* **2002**, *4*, 11–22. [CrossRef]
3. Janssen, M.; Charalabidis, Y.; Zuiderwijk, A. Benefits, adoption barriers and myths of open data and open government. *Inf. Syst. Manag.* **2012**, *29*, 258–268. [CrossRef]
4. Attard, J.; Orlandi, F.; Scerri, S.; Auer, S. A systematic review of open government data initiatives. *Gov. Inf. Q.* **2015**, *32*, 399–418. [CrossRef]
5. Meijer, A.J.; Curtin, D.; Hillebrandt, M. Open government: Connecting vision and voice. *Int. Rev. Adm. Sci.* **2012**, *78*, 10–29. [CrossRef]
6. Jetzek, T.; Avital, M.; Bjørn-Andersen, N. Data-driven innovation through open government data. *J. Theor. Appl. Electron. Commer. Res.* **2014**, *9*, 100–120. [CrossRef]

7. Latre, M.A.; López-Pellicer, F.J.; Nogueras-Iso, J.; Béjar, R.; Zarazaga-Soria, F.J.; Muro-Medrano, P.R. Spatial Data Infrastructures for environmental e-government services: The case of water abstractions authorisa-tions. *Environ. Model. Softw.* **2013**, *48*, 81–92. [CrossRef]
8. Lacroix, P.; Moser, F.; Benvenuti, A.; Piller, T.; Jensen, D.; Petersen, I.; Planque, M.; Ray, N. MapX: An open geospatial platform to manage, analyze and visualize data on natural resources and the environment. *SoftwareX* **2019**, *9*, 77–84. [CrossRef]
9. Ran, J.; Nedovic-Budic, Z. Online decision support infrastructures for integrating spatial planning and flood risk management policies. *ISPRS Int. J. Geo-Inf.* **2024**, *13*, 53. [CrossRef]
10. Jacob-Loyola, N.; Muñoz-La Rivera, F.; Herrera, R.F.; Atencio, E. Unmanned Aerial Vehicles (UAVs) for physical progress monitoring of construction. *Sensors* **2021**, *21*, 4227. [CrossRef] [PubMed]
11. de Souza, A.D.; de Souza, J.S.C.; Rocha, J.G.L.; Sobral, M.C.; Candeias, A.L.B. Using GIS to monitor public works. *Int. J. Gov. Audit.* **2016**, *43*, 30–33. Available online: https://intosaijournal.org/journal-entry/using-gis-to-monitor-public-works/ (accessed on 11 August 2026).
12. International Organization of Supreme Audit Institutions (INTOSAI). *ISSAI 5540: Use of Geospatial In-formation in Auditing Disaster Management and Disaster-Related Aid*; INTOSAI: Vienna, Austria, 2013. Available online: https://www.issai.org/pronouncements/use-of-geospatial-information-in-auditing-disaster-management-and-disaster-related-aid/ (accessed on 11 August 2026).
13. Open Geospatial Consortium (OGC). OGC Standards and Supporting Documents. Available online: https://www.ogc.org/ (accessed on 11 August 2026).
14. Goodchild, M.F.; Fu, P.; Rich, P. Geographic information sharing: The case of the Geospatial One-Stop por-tal. *Ann. Assoc. Am. Geogr.* **2007**, *97*, 250–266. [CrossRef]
15. Morales Escobar, A.A.; Medina Daza, R.J.; Páez Lancheros, A. Medición del desempeño del gobierno de la Infraestructura Colombiana de Datos Espaciales en la administración de tierras con un modelo balance scorecard. *Tecnura* **2024**, *28*, 12–38. [CrossRef]
16. Flyvbjerg, B. What you should know about megaprojects and why: An overview. *Proj. Manag. J.* **2014**, *45*, 6–19. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.