

DPGs for Climate Action Collection



Framework Compliance Summary:

All DPGs included in the DPGs for Climate Action Collection must meet the requirements outlined in the accompanying [identification framework](#).

To demonstrate compliance, each applying DPG is required to complete a structured submission. The information below reflects the self-reported responses provided as part of that process.

GeoPrism Registry

DPG Type: Software

DPG Compliance & Profile Page: <https://www.digitalpublicgoods.net/r/geoprism-registry>

Description: GeoPrism Registry allows for the simultaneous hosting, management, regular update, and sharing of master lists, associated hierarchies, and geospatial data for the geographic objects core to sustainable development including public health, education, agriculture, infrastructure, and other public and private services.

Climate Alignment Assessment

Climate-Relevant Problem it Addresses: Climate action is constrained by fragmented, inconsistent, and non-interoperable geospatial data across sectors such as environment, health, infrastructure, and disaster management. Critical datasets describing climate hazards, vulnerable populations, hydrography, and infrastructure are maintained by different authorities using incompatible identifiers, classifications, and spatial representations. This prevents the integration of data into coherent, analysis-ready systems needed for mitigation, adaptation, resilience planning, and loss and damage assessment.

As a result, decision-makers cannot easily answer fundamental questions such as which populations are at risk from climate hazards, how infrastructure dependencies amplify impacts, or how interventions should be prioritized geographically.

GeoPrism Registry addresses this by establishing a Common Geo-Registry (CGR) that provides a single, authoritative, and temporally-aware source of truth for geospatial entities—ensuring consistent identity, classification, geometry, and relationships across domains. This foundation is essential for enabling LLM-based GeoAI decision support tools, which require structured, semantically consistent, and interconnected spatial data to reason over geographic systems. By normalizing geospatial knowledge into interoperable spatial knowledge graphs, GPR makes it possible for non-technical users to ask complex, location-based questions in natural language and receive reliable, context-aware insights for climate decision-making.

Climate-Relevant Objectives:

- Mitigation
- Adaptation and resilience
- Loss and damage
- Cross-cutting (data infrastructure, interoperability, digital MRV)

Sector Applications:

- Agriculture, forestry and other land use
- Building
- Energy
- Transport
- Waste

Common Deployment Settings

- Regional
- National

International Frameworks it Aligns With

- UN Integrated Geospatial Information (IGIF)

Climate Outcomes Facilitated

Measurable Outcomes it Can Help Facilitate:

- Improved data quality and reduced uncertainty
- Strengthened adaptation planning and resilience capacity
- Enhanced inter-agency coordination and governance

Specific Decisions or Actions it Has Helped Facilitate: GeoPrism Registry enables data-driven, location-based decision-making across climate mitigation, adaptation, and resilience by integrating infrastructure, environmental, population, and health data into a unified spatial knowledge framework.

The U.S. Army Corps of Engineers (USACE) uses the solution to assess the cascading and downstream impacts of critical infrastructure failure, such as levee breaches and reservoir inundation scenarios. This includes identifying affected populations, schools, and essential services, enabling more effective emergency preparedness and response planning. It also supports prioritization of mitigation investments by revealing where infrastructure interventions will have the greatest impact, while reducing duplication of effort across geographically overlapping projects.

Through the Open Geospatial Consortium (OGC) Climate and Disaster Resilience Pilot (CDRP), requirements from Natural Resources Canada (NRCan) and the Emergency Management Organization in Manitoba highlight its role in disaster response planning—particularly in understanding how flooding events disrupt transportation networks, affecting evacuation routes and supply chain logistics for delivering critical resources into impacted areas.

In public health contexts, the solution supports climate-informed interventions by enabling analysis of extreme heat exposure, proximity to water bodies associated with vector-borne diseases, and accessibility of health services via transportation networks. This allows health authorities to target vulnerable populations, optimize resource allocation, and coordinate response strategies that account for both environmental and infrastructure constraints.

More broadly, GeoPrism Registry enables cross-sector institutions—such as ministries of environment, health, and infrastructure; emergency management agencies; and development organizations—to make coordinated, evidence-based decisions grounded in a shared understanding of geographic context.

Deployment Countries: Laos, Mozambique, United States of America

Adopting Organisations: MOH in Laos, the National Geospatial Agency of Mozambique (ADE), the U.S. Army Corps of Engineers (USACE) Civil Works Division, and the U.S. Department of the Interior (DOI)

Key Performance or Impact Indicators: GeoPrism Registry enables measurable improvements in climate decision-making, resilience, and data interoperability by transforming fragmented geospatial data into integrated spatial knowledge systems accessible via map and chatbot interfaces.

Key indicators include:

- Reduction in analysis time for complex climate questions (from days/weeks to minutes) through natural language GeoAI queries
- Number of non-technical users (e.g., planners, emergency managers, public health officials) using the system for decision support
- Improved identification of cascading impacts, including populations and infrastructure indirectly affected through networks (e.g., disrupted transportation or downstream flooding)
- Reduction in time to generate a common operating picture during disaster events
- Increase in interoperable datasets integrated across sectors using consistent geospatial identifiers (CGR)
- Reduction in duplicate or inconsistent geospatial records, improving data quality and trust (IPT)
- Number of institutions and use cases leveraging outputs for resilience planning, emergency response, and climate-health interventions

These indicators demonstrate how GeoPrism expands access to actionable climate intelligence while improving the speed, quality, and coordination of decision-making.

Technical Assessment

Modular Design: GPR is designed using a Model-Driven Engineering (MDE) architecture pattern, enabling a high degree of modularity, configurability, and reuse.

At its core, GPR employs a metamodel abstraction that defines geospatial object classes, attributes, relationships (edge classes), and spatial data structures at runtime—effectively functioning as a geo-ontology schema for spatial knowledge graphs (SKGs). This allows domain models (e.g., water infrastructure, populated areas, health facilities, transportation networks, and flood scenarios) to be defined declaratively rather than programmatically, eliminating the need for custom code when adapting the system to different countries or organizational contexts.

The platform follows principles of encapsulation and decoupling by separating the metamodel (schema), data instances, and processing logic, and exposes functionality through standardized interfaces and APIs. GPR is aligned with a spatial knowledge mesh architecture, an extension of the data mesh paradigm, where domain-specific SKGs are managed as independent yet interoperable modules. This approach enables each domain to maintain authoritative control over its data while supporting cross-domain integration through shared spatial and semantic references.

Additionally, GPR manages temporal and dependency relationships across interlinked SKGs, allowing changes in one domain (e.g., infrastructure disruption) to propagate through the graph in a controlled and traceable manner. The resulting modular SKGs and geo-ontology can be composed into Geo-GraphRAG pipelines, enabling LLMs to query and reason over distributed, evolving geospatial knowledge. This architecture ensures that GPR is inherently extensible, reusable, and adaptable to diverse climate and resilience use cases.

Data Processing:

Processing Type	Capabilities
Ingestion	GPR supports multiple ingestion pathways to accommodate diverse data sources and formats. Data can be imported via common geospatial and tabular formats, including Shapefiles, Excel, CSV, and GeoJSON through RESTful APIs. This flexibility enables integration with external systems such as government registries, sensor feeds, satellite-derived datasets, and other authoritative data providers.
Validation	GPR employs a schema-driven validation approach based on its metamodel and geo-ontology framework. Each participating organization operates within a multi-tenant architecture, where it maintains authoritative control over its data and corresponding geo-ontology definitions. These geo-ontologies act as modular validation gatekeepers, enforcing constraints on attribute values, hierarchical structures, and semantic relationships. This ensures that

	all ingested data is consistent, semantically valid, and aligned with domain-specific standards before being processed or published.
Dissemination	GPR supports flexible data dissemination through both APIs and exportable data products. Because the platform manages temporal versioning and change over time—including updates to attributes, geometries, classifications, relationships, and events such as splits and merges—users can generate time-bound snapshots of data for specific periods of validity. Outputs are available in multiple formats, including geospatial data tables, spreadsheets, and spatial knowledge graphs, and can be accessed via RESTful APIs for integration into downstream systems, analytics platforms, or visualization dashboards.

Data Extraction Mechanism(s): API, GeoJSON, JSON, XLSX

Interoperability, Integration, and Adapters*: *Level 3 Standardized* - Adheres to domain-specific standards, but requires some configuration to link.

GPR is designed using a Model-Driven Engineering (MDE) architecture that is inherently domain-agnostic, making it highly adaptable for integration with external systems, including climate-focused platforms such as MRV systems and NDC dashboards. Rather than being purpose-built for climate applications, GPR enables the declarative definition of geo-ontologies—including geospatial object classes, attributes, and relationships—allowing it to be configured to support a wide range of domain-specific use cases without custom code.

While integration with MRV and NDC systems has not yet been implemented, the architecture is well-suited to these contexts. Climate-related datasets (e.g., emissions inventories, hazard models, or adaptation indicators) can be aligned to consistent geographic references, such as Discrete Global Grid Systems (DGGS) zones or administrative units, and integrated with complementary datasets. For example, population or public health data could be combined with climate indicators to support more comprehensive assessments of exposure, vulnerability, and co-benefits.

GPR exposes data through standards-based RESTful APIs, providing access to geo-objects, relationships, and time-versioned snapshots, enabling external systems to retrieve and reuse authoritative spatial data structures. This approach supports consistent aggregation, comparison across reporting periods, and cross-sector analysis, positioning GPR as a flexible foundation for interoperable climate information systems.

Adoption Readiness:** *Level 3 Portable/Containerized* - The solution is "packaged." It uses industry-standard wrappers (e.g., Docker, Parquet, or Standardized Schemas) that allow it to run or be read in any standard environment with a single command or import.

Scalability, performance, and reliability:

In the OGC CDRP and AI-DGGS pilots, participating organizations demonstrated that DGGS provides a scalable mechanism for integrating and querying geospatial data in near real time, including through natural language interfaces that improve discoverability and accessibility for users. While GPR does not generate DGGS data from source datasets, it complements DGGS-based approaches by managing spatial knowledge graphs that model semantic networks of features needed to understand downstream impacts, cumulative effects, and common operating pictures in climate and disaster resilience scenarios.

GPR architecture supports scalability through modularity, declarative configuration, and standards-based APIs, allowing data and responsibilities to remain distributed across authoritative sources rather than centralized in a single system. For semantic querying, GPR exports spatial knowledge graphs to RDF triple stores, which are designed for large-scale graph querying and can support substantial concurrent demand when deployed with appropriate indexing, replication, and infrastructure scaling. Reliability and disaster recovery depend on the operational environment, but the architecture supports versioned exports, temporal snapshots, backup workflows, and redeployment of graph stores and APIs, which together provide a strong foundation for resilient production deployments.

Computing power: GeoPrism Registry (GPR) can be deployed using standard cloud or on-premise infrastructure typical for web applications, APIs, and graph databases, without requiring specialized high-performance computing. Resource usage scales with data volume and query complexity, particularly for spatial knowledge graph operations, and energy consumption is therefore primarily determined by the deployment environment rather than any uniquely intensive processing within GPR itself.

While GPR is not specifically optimized for low-bandwidth environments, its Model-Driven Engineering (MDE) architecture supports efficient operation in lower-resourced settings by enabling geo-ontologies and data models to be defined declaratively rather than through custom code. This makes it well-suited for environments where technical capacity is limited, while still enabling advanced capabilities such as Geo-GraphRAG workflows to support on-demand, decision-oriented geospatial analysis.

Adoption Evidence and Additional Resources

The deployment of GeoPrism for USACE is highlighted in the [Open Geospatial Consortium Climate and Disaster Resilience Pilot for 2024](#), which includes an instance with USACE data is listed: <https://genai.usace.geoprism.net/>

TerraFrame's work with NRCan and the Emergency Management Organization for the province of Manitoba, Canada, is described in the [OGC AI-DGGS Pilot Engineering Report](#) which includes a link to an instance with data from Canada is listed: <https://ai-dggs.geoprism.net>

GeoPrism is also listed in [The Global Goods Guidebook – Climate Services for Health Annex](#) from Digital Square

* Interoperability, Integration, and Adapters are evaluated in 5 levels:

Level 1 Isolated: Uses proprietary formats or hardcoded logic; requires custom "glue code" or manual conversion to work with external climate tools.

Level 2 Compatible: Data/software can be exported or integrated using common formats (e.g., CSV, JSON), but lacks automated synchronization or shared metadata.

Level 3 Standardized: Adheres to domain-specific standards (e.g., NetCDF/HDF5 for data, OGC APIs for software) but requires some configuration to link.

Level 4 Integrated: Uses machine-readable schemas and versioned APIs; external systems can "plug in" and pull climate variables without manual intervention.

Level 5 Ecosystem-Ready: Fully modular; follows "FAIR" principles (Findable, Accessible, Interoperable, Reusable) and supports automated cross-platform workflows (e.g., a climate model automatically pulling from your dataset).

*Adoption readiness perspective is assessed in 5 levels:

Level 1 Experimental/Raw: The solution exists as "source material" only. It requires significant manual effort, custom scripts, or compilation to become functional. There is no automated setup, and the user must "build" the environment from scratch.

Level 2 Documented/Structured: The solution is organized and includes instructions. Requirements and dependencies are clearly listed, but the setup process is still manual.

Level 3 Portable / Containerized: The solution is "packaged." It uses industry-standard wrappers (e.g., Docker, Parquet, or Standardized Schemas) that allow it to run or be read in any standard environment with a single command or import.

Level 4 Orchestrated / Cloud-Optimized: The solution is designed for modern infrastructure. It includes "Infrastructure as Code" (e.g., Terraform, Helm, or Crawlable Data Catalogs) that allows for automated deployment into cloud environments (AWS/Azure/GCP) with built-in scaling and management.

Level 5 Productized / Plug-and-Play: It offers a zero-friction experience, such as a managed API, a Serverless function, or a Public Data Marketplace listing. A user can gain value or insights within minutes without managing any underlying infrastructure.