Network Working Group   G. S. Kartha
Internet-Draft   Independent
Intended status: Informational   20 January 2026
Expires: 24 July 2026

Geospatial Resource Discovery (GRD): Problem Statement and Conceptual Architecture

draft-kartha-grd-01

Abstract

The Internet provides standardized mechanisms for resolving identifiers, such as domain names, to network locations. These mechanisms assume that a client already knows the identifier of the resource it wishes to access. Emerging classes of applications, including augmented reality, autonomous systems, robotics, and spatial computing, operate under a different assumption. In these systems, discovery begins with physical presence and observer context rather than with a predefined name.

This document describes the architectural gap addressed by a Geospatial Resource Discovery service. It defines the problem space, explains why existing standards are insufficient, and outlines a conceptual architecture for spatial discovery. This document is informational and architectural in nature. It does not define a protocol, data format, or implementation.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Internet architecture has historically focused on resolving semantic identifiers into network endpoints. Systems such as the Domain Name System enable scalable communication by assuming that clients already know the name of the resource they are attempting to reach.

Emerging spatial systems invert this assumption. In augmented reality, autonomous navigation, and smart environments, discovery begins with physical location and context. A user or machine operating at a specific coordinate needs to ask what digital information is associated with the surrounding space and visible volume.

At present, this type of discovery occurs inside closed ecosystems. Each platform maintains its own spatial maps, registries, and access mechanisms. While effective within individual systems, these approaches do not interoperate. This document argues that spatial discovery represents a missing architectural layer and describes the requirements for an open Geospatial Resource Discovery capability.

This document is intended to inform future work on interoperable spatial discovery mechanisms, potentially including protocol specifications.

2. Scope and Non-Goals

2.1. Scope

This document focuses on the application layer. It describes a discovery pattern that can operate over existing transports such as HTTP/3 or CoAP. It intentionally avoids assumptions about rendering engines, device capabilities, or user interface models.

2.2. Non-Goals

This document does not propose changes to IP routing or DNS semantics. It does not define legal ownership of physical space, geometric rendering standards, or a global governance model for spatial authority.

3. Terminology

Spatial Resource: A digital asset associated with a physical location or volume, such as an augmented reality annotation, navigation marker, or machine-readable instruction.

Spatial Cell: A discrete geospatial region used for indexing and lookup, typically derived from a hierarchical grid system.

Frustum: A three-dimensional volume defined by an observer’s position, orientation, and field of view.

Trust Anchor: A cryptographic root of trust (e.g., a Public Key) configured on the client and used to validate data associated with a specific Semantic Layer (e.g., "Safety" or "Transit") within a given context.

4. Problem Statement

The Internet lacks a standardized mechanism to discover digital resources based on physical location and observer context.

Existing discovery mechanisms resolve names rather than locations. While geographic identifiers exist (e.g., [RFC5870]), there is no standard protocol that allows a client to query a location and receive a list of available spatial resources.

Spatial data is also highly application-specific. Autonomous vehicles, wearable devices, and smart infrastructure systems operating in the same physical environment cannot discover or consume each other's data unless they share a proprietary backend.

In addition, current application protocols do not natively support volumetric queries. Requests such as retrieving resources within a three-dimensional viewing volume are not commonly supported in interoperable discovery mechanisms.

Finally, many spatial resources are dynamic and time-sensitive. Real-time spatial information, such as hazards or temporary instructions, requires low latency and short validity periods. Traditional caching assumptions are poorly suited to these constraints.

5. Design Guidelines

Any future specification for Geospatial Resource Discovery should follow a small set of architectural principles. The design should remain independent of transport and operate over existing Internet protocols. Complex geometric calculations should be performed by the client to preserve resolver scalability. Discovery should support semantic filtering so that clients can request only relevant categories of information. Privacy considerations need to be integral to the design and should limit the precision of location data exposed to discovery services.

6. Conceptual Architecture

This section outlines a conceptual architecture referred to as Geospatial Resource Discovery (GRD).

The physical world is divided into discrete spatial cells using a standard hierarchical indexing system. Representing space in this way allows spatial discovery to be reduced to key-based lookups rather than continuous geospatial computation.

```
          +---------+---------+---------+
          | Cell A  | Cell B  | Cell C  |
          |         |         |         |
          +---------+---------+---------+
          | Cell D  | Cell E  | Cell F  |
          | (Index) | (Index) | (Index) |
          +---------+---------+---------+
          | Cell G  | Cell H  | Cell I  |
          |         |         |         |
          +---------+---------+---------+
```

Figure 1: Hierarchical Spatial Grid

A GRD resolver functions as a directory service. The term "resolver" is used descriptively and does not imply a standardized network role. It maintains references to spatial resources indexed by cell identifier and returns pointers to the actual assets. The resolver does not host large content objects such as three-dimensional models.

Authority over physical space is not strictly hierarchical. Multiple independent entities may publish resources associated with the same spatial cell. This introduces a fundamental architectural choice between client-side aggregation of multiple authorities and resolver-side federation of records.

For example, a federated resolver simplifies client logic and can vouch for aggregated data integrity, but requires coordination among authorities. Client-side aggregation offers greater autonomy and avoids a central point of coordination but places the burden of trust evaluation and result merging on each client.

7. Volumetric Query Concept

To balance scalability and privacy, GRD follows a coarse query and fine filter model.

The client computes its precise view frustum locally and determines which standard spatial cells intersect that volume. It then queries the resolver for resources associated with those cells, optionally filtered by semantic category. This translation from a continuous geometric volume to discrete cell identifiers allows spatial discovery to leverage standard, cache-friendly HTTP request semantics for what is inherently a spatial query.

```
                   / \
                  /   \  <-- View Frustum
                 /     \
                /       \
         +-----/---------/-----+
         | Cell X | Cell Y |   |  <-- Client requests Cell X + Y
         +--------+--------+---+      (Standard HTTP GET)
```

Figure 2: Frustum-to-Cell Intersection

The client performs final filtering locally and discards resources that fall outside the exact viewing volume. This approach ensures that the resolver does not learn the client’s precise orientation or gaze direction.

8. Relationship to Existing Work

8.1. IETF Standards

RFC 5870 (’geo’ URI) defines a method to identify a physical location but does not define a mechanism to retrieve resources associated with that location. GRD complements RFC 5870 by providing the resolution layer.

RFC 9176 (CoRE Resource Directory) enables resource discovery on constrained networks but relies on network endpoints rather than geodetic context. It does not support volumetric or spatial-context queries (e.g., "what is visible from this coordinate").

RFC 6763 (DNS-SD) is effective for local service discovery via multicast or administrative domains but cannot scale to global, location-based queries required for open spatial computing.

8.2. OGC and GIS Standards

OGC standards (e.g., WFS, OGC API - Features) are designed for complex querying of heavy geospatial databases. GRD is designed as a lightweight, low-latency resolution service (conceptually similar to DNS in its indirection role) for real-time spatial computing, handing off the heavy data transfer to application-specific protocols.

9. Security and Trust Challenges

A primary risk in spatial discovery is reality poisoning: the injection of false or malicious information that directly alters a user’s perception of the physical world. The safety implications are significant--consider the effect of a counterfeit right-of-way marker for an autonomous vehicle or a hallucinated safe path over a physical hazard for a first responder.

9.1. Trust Model

GRD cannot rely on a single global root of trust. Instead, trust must be contextual and scoped by semantic layer. A "Safety" layer might rely on a government Trust Anchor, while a "Commercial" layer might rely on a different registry. The architecture is expected to require that Spatial Resource Records (SRRs) carry cryptographic signatures verifiable against the client’s configured Trust Anchors for that specific layer.

10. Privacy Considerations

10.1. Attacker Model

The GRD Resolver is modeled as an Honest-but-Curious adversary. It is trusted to return correct spatial records but is assumed to log queries in an attempt to reconstruct client trajectories.

10.2. Measurable Privacy Goals

Privacy risks are mitigated by enforcing minimum cell granularity (k-anonymity), intended to ensure that a query does not uniquely identify a user within a small physical area. Additionally, the architecture is expected to avoid reliance on persistent session identifiers that allow the Resolver to correlate separate queries to a single device (Unlinkability).

11. IANA Considerations

This document does not request any IANA actions.

12. References

[RFC5870] Mayrhofer, A. and C. Spanring, "A Uniform Resource Identifier for Geographic Locations (’geo’ URI)", RFC 5870, June 2010, <https://www.rfc-editor.org/rfc/rfc5870>.

[RFC9176] Shelby, Z., Koster, M., and C. Bormann, "Constrained RESTful Environments (CoRE) Resource Directory", RFC 9176, April 2022, <https://www.rfc-editor.org/rfc/rfc9176>.

[RFC6763] Cheshire, S. and M. Krochmal, "DNS-Based Service Discovery", RFC 6763, February 2013, <https://www.rfc-editor.org/rfc/rfc6763>.

Appendix A. Illustrative Data Model

This appendix is non-normative and illustrates the type of data a resolution service might return.


```
{
  "cell_id": "8928308280fffff",
  "resources": [
    {
      "id": "res-12345",
      "type": "safety.warning",
      "location": { "lat": 52.3, "lon": 4.9, "alt": 10 },
      "layer": "infrastructure",
      "uri": "https://city-assets.example.com/signs/stop.glb",
      "validity_ms": 60000,
      "signature": "ab382...[cryptographic proof]"
    }
  ]
}
```


Figure 3

Author's Address

Gokul Sivasankaran Kartha
Independent
Email: kartha.gokul@gmail.com