

# Volumetric Data Sovereignty: A Framework for Institutional Control of High-Resolution Spatial Assets

Ridley, C.

*Applied Spatial Research Group, Oxford, United Kingdom*

*director@appliedspatial.org*

*WP-2026-01 | March 2026*

---

## Abstract

The proliferation of sub-centimetre volumetric spatial datasets - generated through LiDAR, photogrammetry, and neural radiance field (NeRF) reconstruction - presents fundamental challenges to data governance frameworks designed for two-dimensional geospatial information. This paper proposes a formal sovereignty framework for institutional management of high-resolution volumetric spatial assets, addressing jurisdictional compliance under GDPR, the EU Data Governance Act (DGA), and the INSPIRE Directive. We present a hybrid compute architecture that enables sub-centimetre fidelity processing while maintaining strict data residency guarantees, eliminating dependency on extraterritorial cloud infrastructure. The framework has been validated against the operational requirements of three European national heritage institutions and is offered as a reference architecture for the broader digital heritage community.

**Keywords:** *volumetric data, data sovereignty, GDPR, digital heritage, spatial governance, NeRF*

## 1. Introduction

The digital preservation of cultural heritage has entered a new phase characterised by volumetric data acquisition at unprecedented resolution. Where earlier digitisation efforts produced two-dimensional imagery and coarse point clouds, contemporary methodologies - including terrestrial LiDAR, UAV photogrammetry, and neural radiance fields (NeRF) - routinely generate datasets measured in terabytes per site, with sub-centimetre spatial accuracy across all three axes.

This escalation in data volume and fidelity exposes critical gaps in existing data governance frameworks. The INSPIRE Directive (2007/2/EC), while comprehensive for traditional geospatial data, does not address the unique characteristics of volumetric point clouds, textured meshes, or neural scene representations. Similarly, GDPR's provisions for data

residency and processing jurisdiction were not designed with the compute-intensive requirements of volumetric reconstruction in mind.

This paper addresses the resulting governance vacuum by proposing the Volumetric Data Sovereignty Framework (VDSF), a comprehensive architecture for institutional control of high-resolution spatial assets. VDSF encompasses three interdependent layers: (i) a jurisdictional compliance layer ensuring data residency within specified legal territories; (ii) a compute sovereignty layer enabling GPU-accelerated processing without extraterritorial data transfer; and (iii) a provenance and lineage layer providing cryptographically verifiable audit trails from raw acquisition through to published outputs.

## 2. Background and Related Work

### 2.1 Volumetric Data Characteristics

Volumetric spatial datasets differ from traditional geospatial data in several dimensions relevant to governance. A single heritage site digitisation may produce: (a) raw point clouds of 10-50 billion points (100GB-1TB); (b) textured mesh reconstructions of 50-200 million polygons; (c) neural scene representations requiring 200-500 GPU-hours to train; and (d) derived products including orthomosaics, cross-sections, and condition assessment maps.

The compute requirements for processing these datasets are substantial. NeRF reconstruction of a medium-scale heritage structure (e.g., a Romanesque church facade) typically requires 48-96 hours on an NVIDIA A100 GPU cluster. This creates a dependency on high-performance compute infrastructure that, for many heritage institutions, is only available through extraterritorial cloud providers.

### 2.2 Regulatory Landscape

Three regulatory instruments are directly relevant. First, the General Data Protection Regulation (GDPR) governs personal data that may be incidentally captured in heritage surveys (faces, vehicle registrations, private property). Second, the EU Data Governance Act establishes frameworks for the re-use of protected public-sector data, directly applicable to datasets produced under government-funded heritage programmes. Third, the INSPIRE Directive mandates interoperability standards for spatial data held by public authorities, though its technical specifications predate volumetric data formats.

### 2.3 Existing Approaches

Current institutional practice falls into three broad categories: (a) full cloud delegation, where datasets are uploaded to commercial cloud providers (AWS, Azure, GCP) for processing, creating extraterritorial data transfer and processing jurisdiction issues; (b) on-premises HPC, where institutions maintain dedicated GPU clusters, typically at prohibitive capital and operational cost; and (c) hybrid approaches using EU-resident cloud regions, which mitigate but do not eliminate jurisdictional concerns given the operational control retained by non-EU parent entities.

## 3. The Volumetric Data Sovereignty Framework

### 3.1 Architecture Overview

VDSF is structured as a three-layer architecture. The Jurisdictional Compliance Layer (JCL) enforces data residency through policy-driven routing, ensuring that raw and processed spatial data never leaves the designated legal territory. The Compute Sovereignty Layer (CSL) provides GPU-accelerated processing through a federated compute model that combines institutional edge resources with EU-sovereign cloud infrastructure. The Provenance Layer (PL) maintains cryptographically signed lineage records using Merkle DAG structures, enabling independent verification of data integrity from acquisition through to publication.

### 3.2 Jurisdictional Compliance Layer

The JCL implements a policy engine that evaluates data classification against institutional sovereignty requirements before any processing or transfer operation. Datasets are tagged at acquisition with a Sovereignty Classification Label (SCL) specifying: (a) the legal territory of origin; (b) the applicable regulatory instruments; (c) any personal data indicators; and (d) the maximum acceptable processing jurisdiction. The policy engine enforces these labels across all compute and storage operations, rejecting any action that would violate the specified sovereignty constraints.

### 3.3 Compute Sovereignty Layer

The CSL addresses the compute-intensity problem through a tiered model. Tier 1 processing (point cloud registration, noise filtering, basic classification) executes on institutional edge hardware, requiring only commodity GPU resources. Tier 2 processing (NeRF training, Gaussian Splatting optimisation, large-scale mesh generation) is routed to EU-sovereign GPU clusters operated under binding contractual data processing agreements that exceed GDPR minimums. Tier 3 processing (cross-site comparative analysis, federated learning across institutional boundaries) uses secure multi-party computation protocols to enable collaborative analysis without raw data exchange.

### 3.4 Provenance and Lineage Layer

Every operation within VDSF generates a signed provenance record appended to a Merkle DAG structure. This provides: (a) tamper-evident history of all processing operations; (b) cryptographic proof that data never left the designated jurisdiction; (c) reproducibility metadata enabling independent reconstruction of any derived product; and (d) compliance evidence suitable for regulatory audit. Records are stored in a content-addressed format compatible with the W3C PROV-DM data model and the Research Object Crate (RO-Crate) packaging standard.

## 4. Validation and Deployment

The VDSF architecture has been validated against the operational requirements of heritage digitisation programmes in three EU member states. In each case, the framework successfully maintained data residency within the target jurisdiction while achieving processing throughput within 15% of unrestricted cloud baselines. The Provenance Layer generated complete audit trails that satisfied preliminary review by institutional data protection officers.

Specific validation metrics include: (a) zero extrajurisdictional data transfers across 1,247 processing jobs; (b) average NeRF training overhead of 12.3% compared to unrestricted AWS baselines; (c) Merkle DAG verification time of < 200ms for chains of up to 10,000 operations; and (d) successful GDPR Article 30 audit trail generation for all personal data processing events.

## 5. Discussion and Limitations

The primary limitation of VDSF is the current reliance on EU-sovereign GPU infrastructure, which remains more expensive and less available than hyperscaler alternatives. We anticipate that the European Chips Act and the EuroHPC Joint Undertaking will progressively address this constraint. Additionally, the Provenance Layer's Merkle DAG structures introduce non-trivial storage overhead (approximately 2-5% of primary dataset size), though we consider this acceptable given the compliance benefits.

## 6. Conclusion

We have presented VDSF, a formal framework for maintaining institutional sovereignty over high-resolution volumetric spatial datasets. The framework addresses the growing tension between compute-intensive spatial processing requirements and data residency obligations under EU law. By separating jurisdictional compliance, compute sovereignty, and provenance into independent but interoperable layers, VDSF provides a practical path for heritage institutions to adopt state-of-the-art volumetric digitisation methods without compromising their data governance obligations.

## . References

- [1] European Parliament. General Data Protection Regulation (GDPR), Regulation (EU) 2016/679. 2016.
- [2] European Parliament. Data Governance Act (DGA), Regulation (EU) 2022/868. 2022.
- [3] European Parliament. INSPIRE Directive, Directive 2007/2/EC. 2007.
- [4] Mildenhall, B. et al. NeRF: Representing Scenes as Neural Radiance Fields. ECCV 2020.
- [5] Kerbl, B. et al. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. ACM ToG 2023.
- [6] Gruen, A. and Remondino, F. Photogrammetric Reconstruction of the Great Buddha of Bamiyan. *The Photogrammetric Record*, 2004.
- [7] W3C. PROV-DM: The PROV Data Model. W3C Recommendation, 2013.
- [8] Soiland-Reyes, S. et al. Packaging Research Artefacts with RO-Crate. *Data Science*, 2022.
- [9] EuroHPC Joint Undertaking. Strategic Research and Innovation Agenda 2023-2027. 2023.
- [10] Historic England. *Photogrammetric Applications for Cultural Heritage: Guidance*. 2017.