REPORT

DATA TO DECISIONS

Rebuilding Trust in Decision-Grade Geospatial Data

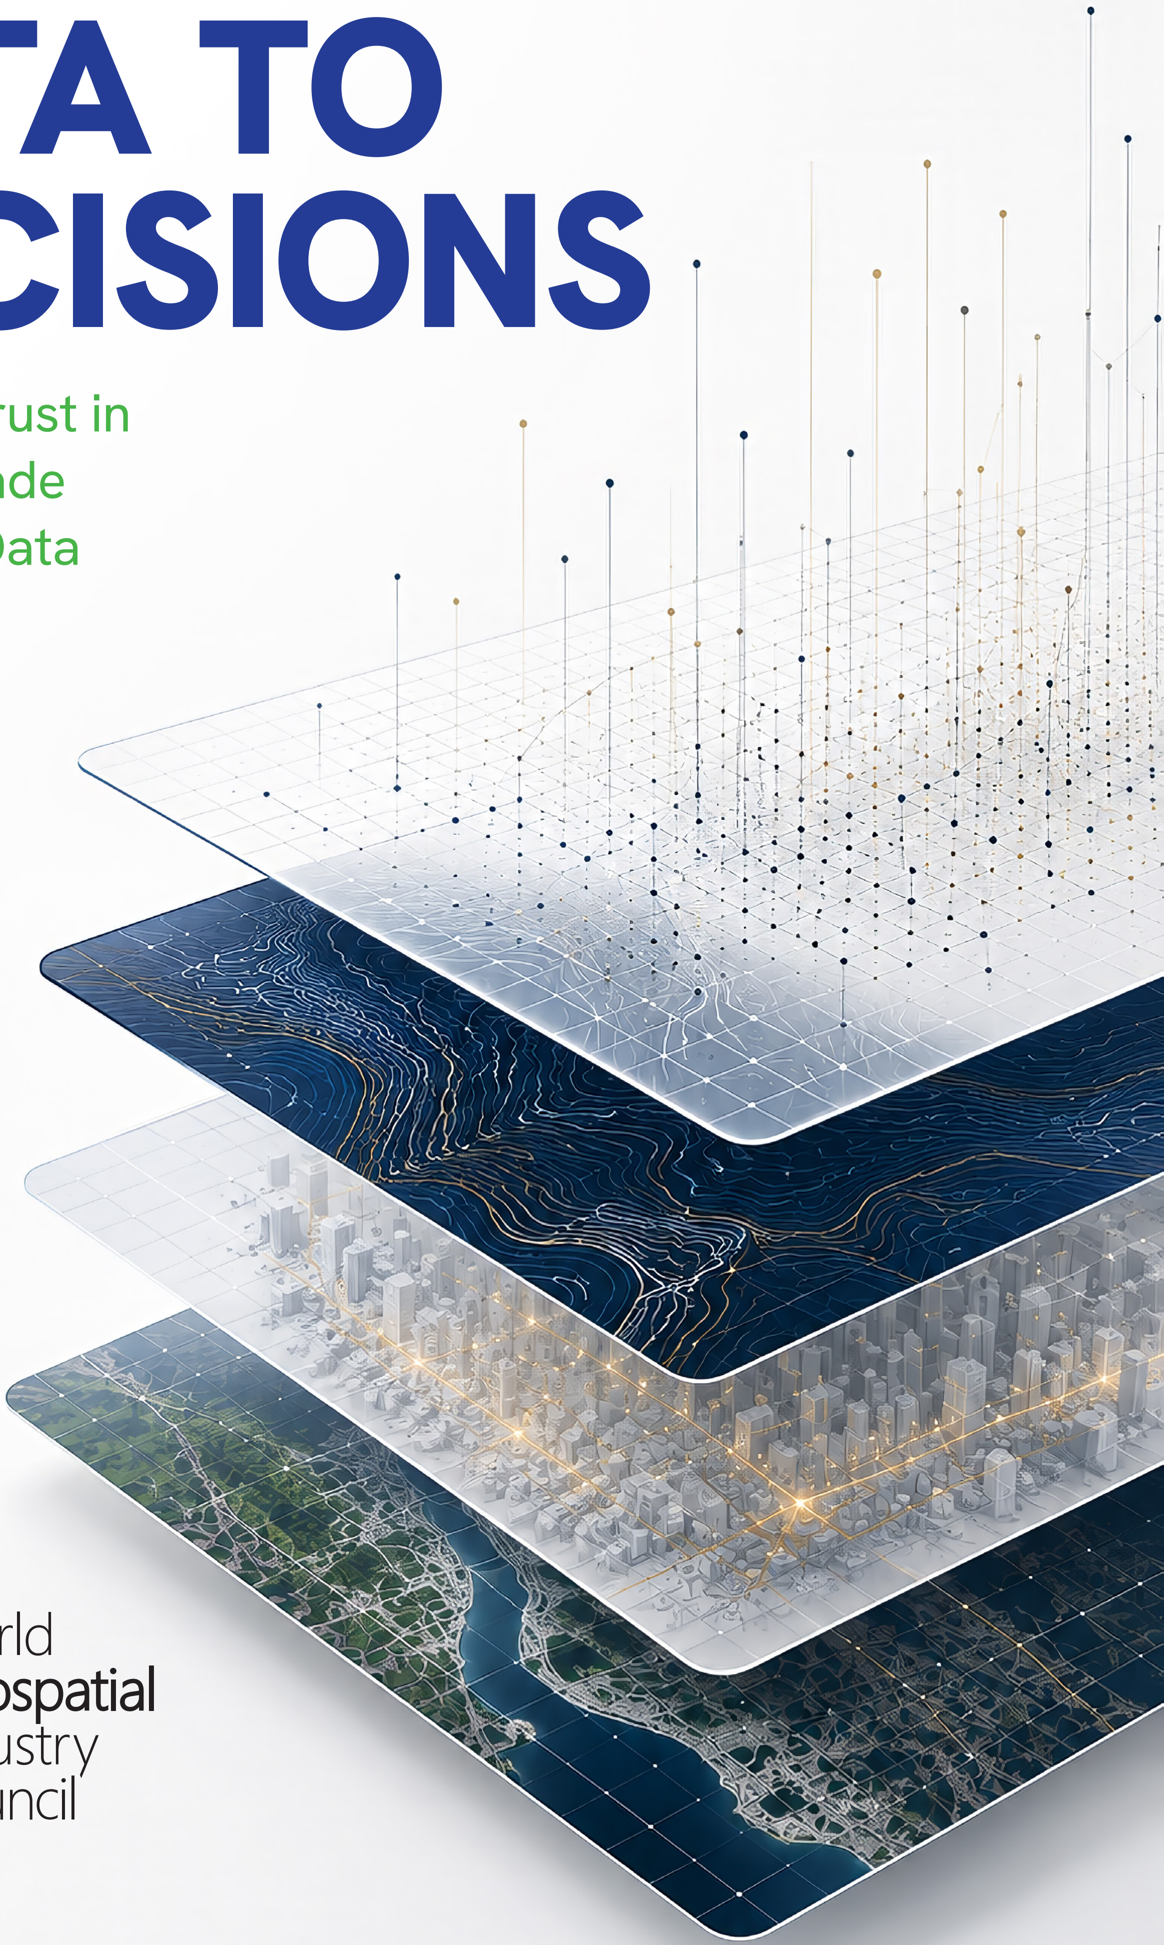

World Geospatial Industry Council

From the Executive Director

The geospatial and EO industry's next phase of growth will not be defined by higher resolution alone. It will be defined by whether we can turn complex, multi-source geospatial inputs into outputs that leaders can trust, defend, and act upon: in short, by whether the data is decision-grade.

That is a harder test than it sounds, and it is not one any single company or organisation can meet on its own. The chain runs the full length of our industry: from the sensor, through the model, to the decision. A weakness at any point travels downstream, so the response has to span the same length. That is why WGIC took on this work. Our members sit at every stage of that chain, and few other bodies can convene the whole of it.

This report is where our effort begins. It gives our members a common reference point at a moment when regulation and AI capability are both moving quickly, and it sets out practical recommendations, not principles alone. Above all, it is not the final word, but an invitation to our members and to the wider industry to build on it with us.

- **Aaron Addison,** Executive Director, WGIC

Foreword

Geographers have long been fortunate in having maps: a way of displaying and analysing disparate and complex datasets. Modern maps may sometimes have distorted the real world for military or political reasons, but we could be confident in the underlying data because we knew how and when it was made. The process was transparent and measurable. The data was not always entirely accurate, but we understood its limitations.

Machine learning and generative AI, coupled with the exponential growth of location-related data, have challenged that confidence and led us to question how far the data can be trusted. How can we be sure we are making critical decisions on decision-grade data when we do not know where it came from, who created it, or what process and technology were used?

This report examines those questions of confidence and trust, and challenges current thinking. Drawing on a range of industry voices and practical expertise, it makes a series of detailed recommendations that will strengthen practice and guide users and practitioners alike through a complex issue. It is timely and essential reading for anyone making critical decisions with geospatial data.

- **John Renard,** Chair, WGIC Data Committee

Executive Summary

The geospatial sector is moving from measurement-based trust, where we believed maps because we knew how they were made, to model-based trust, where data is inferred, synthesised, and generated at scale by AI systems whose workings are hard to inspect. This shift creates enormous value alongside new, largely invisible risks: silent errors, untraceable lineage, and blurred accountability when decisions go wrong.

This is not a challenge unique to geospatial: every sector deploying AI faces the same questions of provenance and accountability, and these expectations are increasingly written into law, regulation, and contract across industries. But the impact may be more consequential for geospatial data, given decades of established practice that must now be retrofitted. And no technology alone, however well designed, can enforce this: metadata, provenance, and governance require a supporting legal framework of clear rights and obligations, not just better tooling.

Definition: decision-grade data
Geospatial data that is demonstrably fit for a specified decision, with traceable origins, documented transformations, appropriate accuracy and currency, communicated uncertainty, lawful usage rights, and identifiable accountability. Fitness is always relative to the decision, use case, jurisdiction, timeframe, and consequence level.

Four foundations of trust

- Provenance
- Governance
- Metadata
- Accountability

Three priority actions

- Embed provenance & machine-readable metadata now
- Convene an industry-wide governance umbrella
- Define contractual liability across the data chain

How to read this report

Executives and policy leads should start with Chapter 6; legal, risk, and compliance teams with Chapters 2 and 4; and data engineers and product teams with Chapters 3 and 5.

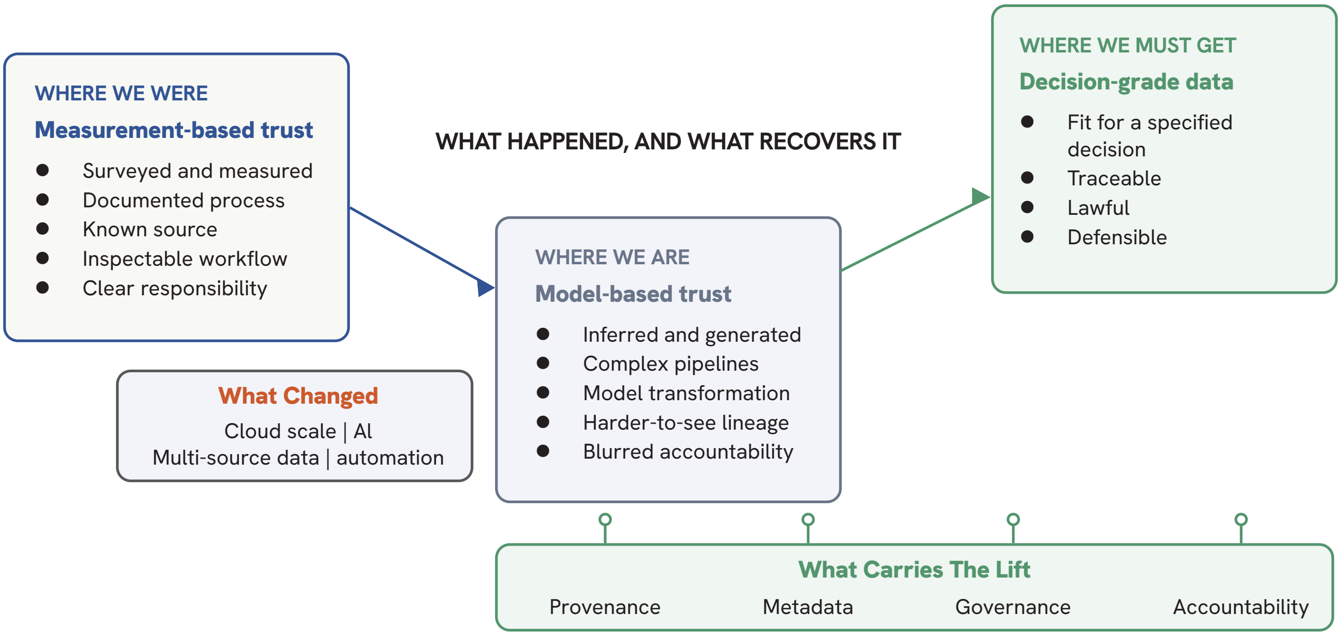

Decision-grade data is not a return to measurement-based trust. It is a different and higher state.

Figure 1. The shift from measurement-based to model-based trust.

This report argues that trust must be rebuilt on four foundations (provenance, governance, metadata, and accountability) delivered through emerging standards, federated architectures, and governance-ready Geospatial Foundation Models. Chapter 2 develops these foundations into six linked properties of decision-grade data – source and capture, integrity and authenticity, provenance and lineage, quality and uncertainty, rights and governance, and fitness for decision – using consistent terminology throughout. WGIC, whose members span every stage of the geospatial value chain, is uniquely placed to convene that effort.

This report is organised around three questions: what has changed, what makes geospatial data decision-grade, and what organisations and WGIC should do next.

WHAT HAS CHANGED?	**Chapter 1: From Measurement to Models** how trust was built, and what cloud and AI changed
WHAT MAKES GEOSPATIAL DATA DECISION-GRADE?	**Chapter 2: What Makes Data Decision-Grade?** the six properties, and why none is sufficient alone
	Chapter 3: Building Trust into the Data Lifecycle metadata, lineage, hashing, and human review
	Chapter 4: Sovereignty, Transparency and Control jurisdiction, licensing, federation and Article 50
	Chapter 5: Governance by Design governance built into the model architecture, not layered on top
WHAT SHOULD ORGANISATIONS AND WGIC DO NEXT?	**Chapter 6: WGIC Action Agenda** who acts, when, and what WGIC should lead

Four of the six chapters answer the same question.

Contents

■ CHAPTER 1

From Measurement to Models

In one sentence:

We trusted maps because we knew how they were made; cloud and AI now generate geospatial data at a scale and speed that make that link harder to see.

How trust was built

For most of the twentieth century, geospatial data was trusted because its origins were understood. National mapping agencies, hydrographic offices, and specialist survey firms operated within well-defined professional and legal frameworks. Data was captured deliberately: surveyors established control points, aircraft flew calibrated missions, and cartographers interpreted imagery by eye. Every step left a trace, and responsibility was clear.

Because capture was costly and labour-intensive, major surveys were conducted periodically rather than continuously. Users worked with authoritative datasets whose age, accuracy and update limitations were generally known. Trust emerged not from novelty or volume, but from repeatability, professional judgment, and institutional memory.

The map was a stable artefact, grounded in physical measurement and documented workflows. That trust was relative, not absolute: traditional mapping could still embed outdated information or institutional bias. What distinguished it was that its production process was inspectable. The same quality architecture principles still apply today, now formalised in standards such as ISO/IEC 25024 and ISO 5725 [1][2] alongside ISO 19157-1 [16], even where a document does not name them individually.

What changed: cloud and AI at scale

Cloud computing and AI have lowered the barriers to creating and distributing geospatial data while changing how it is produced. Cloud platforms can process imagery, LiDAR and sensor data at a scale that was previously impractical, and AI now automates feature extraction and change detection at a pace no human team could match.

Near-real-time flood extent maps, for example, can be produced during extreme weather using automated classification [3]. Two open building-footprint datasets illustrate the scale involved: Microsoft's Global ML Building Footprints, generated by deep learning from satellite imagery, reports approximately 1.4 billion detected buildings derived from imagery acquired between 2014 and 2024 [4], while Google's Open Buildings dataset reports over 1.8 billion building footprints across a 58-million-square-kilometre area covering Africa, South Asia, Southeast Asia, Latin America and the Caribbean [5]. Cloud-native columnar formats such as Parquet and GeoParquet support efficient large-scale analytics and storage, but do not by themselves provide the dataset-level lineage, version governance and provenance controls required for decision-grade use. Organisations therefore need to implement those controls around the format.

This scale comes at a cost: many datasets are no longer the result of deliberate, expert-led cartography, but of opaque pipelines where algorithms make millions of micro-decisions that are difficult to inspect. Data once never consumed on a map (commercial activity, IoT sensor streams) is now being "mapified" at unprecedented volume, extending the reach and the risk of geospatial products beyond the traditional geospatial community.

The flood-map journey

One recurring example threads through this report: a flood map, tracked from measurement through to use. Each stage adds capability, and a new way for trust to break.

THE SAME FLOOD MAP, FIVE STAGES ON

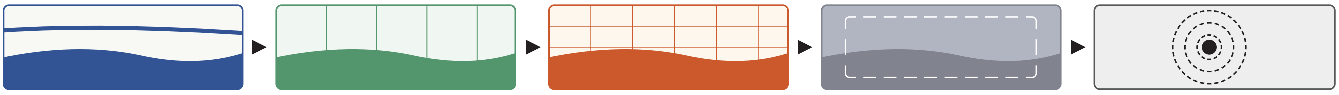

MEASURED	GENERATED	SECURED	GOVERNED	USED
Surveyed and checked against established methods.	Satellite imagery classified by AI and fused with other sources.	Hashed, timestamped, carrying machine-readable lineage.	Data stays in jurisdiction while authorised analysis crosses borders.	Enters a GFM or automated decision process.
WHAT CAN FAIL	**WHAT CAN FAIL**	**WHAT CAN FAIL**	**WHAT CAN FAIL**	**WHAT CAN FAIL**
Slow to produce; may already be ageing.	Temporary floodwater misclassified; versions change silently.	Tampering or substitution goes undetected.	Licensing and jurisdictional limits lost downstream.	Provenance disappears inside the model.
CONTROL REQUIRED	**CONTROL REQUIRED**	**CONTROL REQUIRED**	**CONTROL REQUIRED**	**CONTROL REQUIRED**
Declared accuracy and documented survey method	Uncertainty communication and validation against ground truth	Cryptographic hashing and trusted timestamping	Federated architecture and governance metadata.	Authenticity and sovereignty gates, plus audit trail

Audit trail returns to the decision

Figure 2. The flood-map journey: five stages, and the control that keeps each one trustworthy.

The hierarchy of trust

What is needed is a transition model in which AI-derived outputs are continuously validated against physical evidence, trusted reference data and documented provenance. This provenance-linked reliability creates a hierarchy of trust in modern mapping:

EACH LEVEL RESTS ON THE ONE BELOW	WHAT SECURES IT
4 **DECISIONS** **Logic and provenance:** How the output was used to inform an action.	**Lineage and audit frameworks**
3 **SEMANTICS** **Context and material:** What the data is classified or labelled as.	**Classification metadata**
2 **GEOMETRY** **Spatial integrity:** Coordinate reference systems and positional accuracy.	**Coordinate standards**
1 **HARDWARE** **Physics and calibration:** Sensor state at the moment of capture.	**Calibration metadata**

A failure at a lower level cannot be corrected at a higher one.

"If I can't explain how AI was used at each step of a process to get a certain output, then I can't use it."

WGIC roundtable participant, Geo Week 2026

The traditional signals of trust (named surveyors, declared accuracy, documented workflows) are increasingly absent or hard to observe. And geospatial data is no longer a niche market: through open standards such as the OGC APIs, it is woven into mainstream, non-geospatial products, so the consequences of unclear provenance now reach far beyond the traditional geospatial community.

Recommendations

Producers and platforms
Document where and how AI enters every step of a data production workflow: if it cannot be explained, it cannot be trusted.

Outcome: *Restores an inspectable production process for AI-derived data.*

All members
Adopt the four-level hierarchy of trust (hardware, geometry, semantics, decisions) as an audit framework for modern pipelines.

Outcome: *Gives every pipeline a common language for where trust can break.*

Platforms
Communicate dataset versions and update cycles explicitly, so users know which data informed which decision.

Outcome: *Prevents decisions being made on silently superseded data.*

■ CHAPTER 2

What Makes Data Decision-Grade?

In one sentence:

Decision-grade data rests on six linked properties, none of them sufficient alone.

This chapter sets out the framework that runs through the rest of this report. It expands the four foundations introduced in the Executive Summary – provenance, governance, metadata, and accountability – into six linked, operational properties. Each property below is necessary; none is sufficient alone: a dataset can be authentic and untampered while still being the wrong vintage, the wrong resolution, or unlicensed for the use in question.

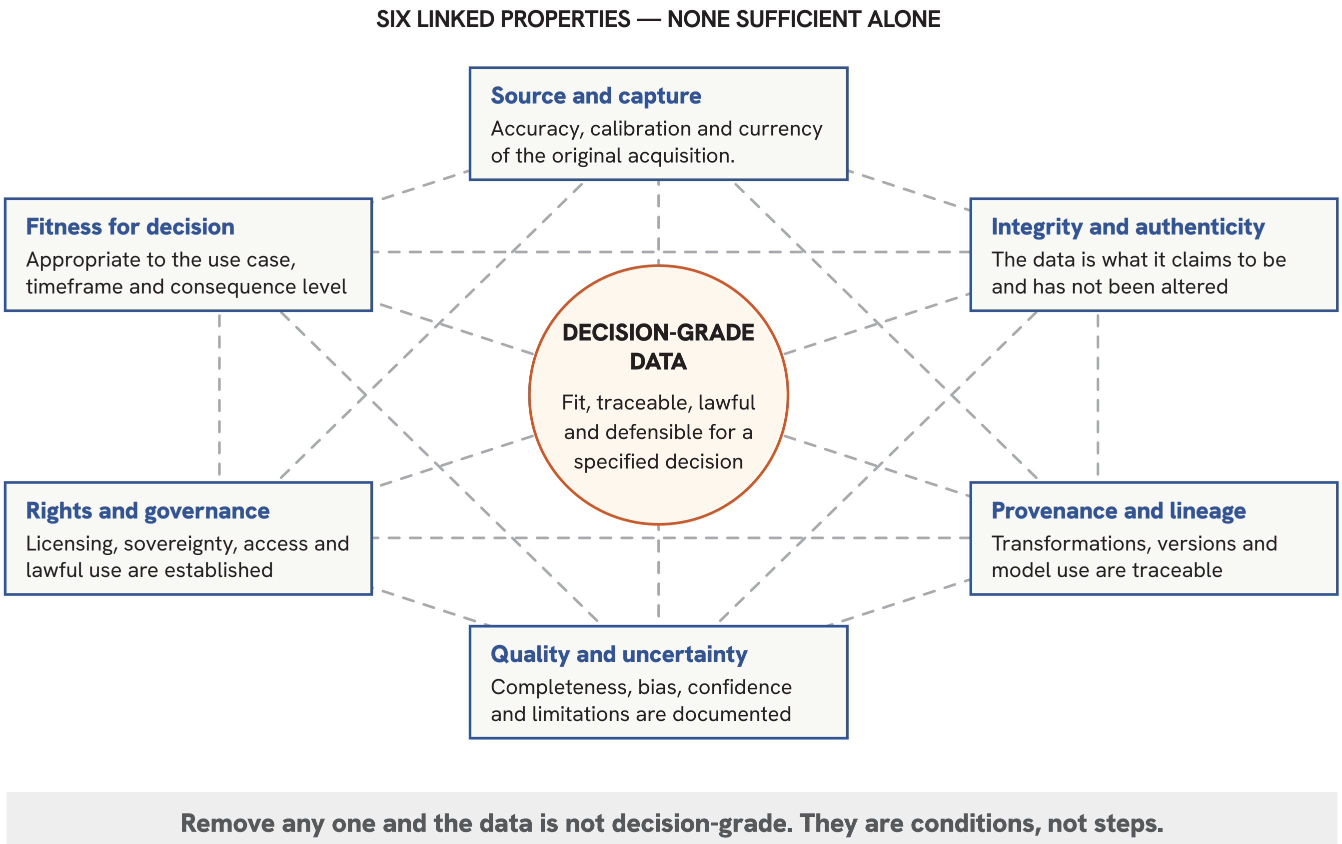

Figure 3. The decision-grade data framework: decision-grade data = fit, traceable, lawful and defensible for a specified decision.

What each control proves, and does not prove

These properties must not be mistaken for one another, or for proof of truth. These controls answer different but related questions: authenticity establishes that data is what it claims to be; integrity, that it has not been altered; provenance, where it came from and how it was transformed. None of them guarantees that the original measurement was accurate, or that the data suits the decision at hand.

Table 1. What trust controls prove, and do not prove.

Control	Proves	Does not prove
Authenticity	The data is what it claims to be.	That it is accurate.
Integrity	It has not been altered.	That the source was correct.
Provenance	Where it came from and how it changed.	That it fits the decision.
Quality	Accuracy, completeness and consistency.	That use is lawful.
Uncertainty	The limits and confidence of the output.	Who bears responsibility.
Fitness for purpose	Suitability for a specified use.	Suitability for every use.

Where this breaks in practice

These issues become acute when poor geospatial data contributes to catastrophic failure, raising hard questions about liability. If an autonomous system crashes because it relied on flawed cloud-processed maps, who is responsible: the data provider, the model developer, the platform operator, or the end user who trusted the output? Traditional regimes assumed clearer lines of authorship, supported by accuracy statements and professional standards; in an AI-driven ecosystem those lines are blurred. Contracts and licensing terms may allocate or limit liability across the value chain, including by transferring some risks downstream.

Data ageing, coverage and human competence

Data quality is time dependent. A dataset that was decision-grade when captured may become progressively less reliable as terrain, infrastructure or population patterns change; metadata should record not only capture date and update cycle, but the conditions under which the data should be reassessed.

Geographic coverage must also be evaluated for bias: densely populated, commercially important areas tend to have richer, more frequently updated data than rural or underserved regions, which can distort AI-derived outputs.

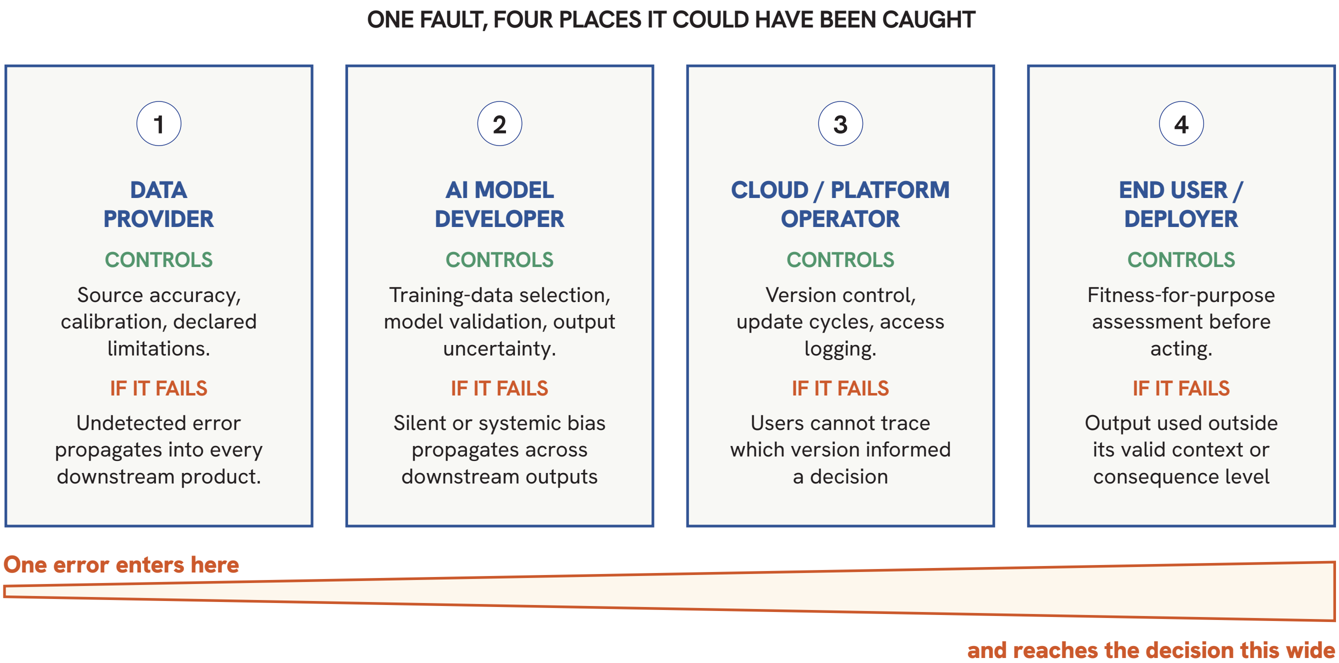

Figure 4. When flawed data causes harm: the control each actor holds, and what happens if it fails.

Illustrative scenario: the phantom road.

A logistics operator relying on automatically updated navigation data routes vehicles onto roads that no longer exist. What failed: no visibility of which dataset version informed the decision. What would have prevented it: version lineage and provenance tracking.

Illustrative scenario: the trap road.

Not every failure is accidental: some maps deliberately include fictitious features (trap streets or trap roads) inserted to detect unauthorised copying. Ingested by an AI pipeline without that context, a deliberate fabrication can be mistaken for real infrastructure. What failed: no mechanism to flag known deliberate anomalies before ingestion. What would have prevented it: provenance metadata that records known trap features alongside genuine ones.

Illustrative scenario: the flood that never receded.

An AI classifier labels temporary floodwater as permanent water bodies, and downstream plans treat it as ground truth. What failed: a visually convincing output masked uncertainty. What would have prevented it: explicit uncertainty communication and validation against ground truth.

As geospatial capabilities become embedded in mainstream analytics, users without specialist training may produce or interpret output without understanding coordinate systems, scale, resolution, or fitness for purpose; governance must address competence and review, not only model performance.

Recommendations

Data producers

Pair every AI-derived product with explicit uncertainty information and rigorous validation for safety-critical uses.

Outcome: *Reduces the risk that visually convincing outputs mask underlying error.*

Legal and procurement

Define responsibility contractually across the chain: provider, model developer, platform, and end user.

Outcome: *Clarifies contractual responsibility when derived data contributes to harm.*

All members

Record data age, update cycle, and reassessment triggers in metadata, and flag known coverage bias by geography.

Outcome: *Lets users judge whether ageing or uneven-coverage data still fits their decision.*

The next three chapters trace trust through three layers of the geospatial data lifecycle: trust in the data itself (Chapter 3), trust in who controls and governs it (Chapter 4), and trust in the models trained on it (Chapter 5).

Building Trust into the Data Lifecycle

In one sentence:

Authenticity becomes infrastructure through machine-readable metadata, verifiable lineage, and cryptographic integrity, validated by people who understand the data.

For users across logistics, real estate, insurance, autonomous systems, and location intelligence, the integrity of geospatial data is directly tied to liability exposure, SLA compliance, and competitive differentiation. A robust authenticity framework must go beyond compliance on report and address the full data supply chain from acquisition to consumption.

Metadata standards

Metadata standards should be evaluated through the lens of machine-readability and pipeline integration. DCAT and GeoDCAT-AP [6][7] enable automated cataloguing and federated discovery across commercial marketplaces, reducing onboarding friction when sourcing third-party datasets.

Schema.org markup [8] can improve the discoverability of geospatial assets within broader enterprise data ecosystems, including cloud-native platforms, though it does not by itself certify quality or licensing.

ISO 19115-1 [9] retains value for regulated procurement, but commercial products increasingly favour standards that integrate natively with REST APIs and modern data-mesh architectures. ISO/IEC 22989 [10] gives the industry a shared vocabulary for AI systems and data, though it does not itself prescribe a geospatial metadata-readiness framework.

National implementations and profiles also matter. In the United States, the FGDC endorses ISO 19115-1 and related ISO geospatial metadata standards while continuing to support its earlier CSDGM legacy standard [11]. Separately, Australia and New Zealand's ANZLIC metadata profile provides a broadly comparable framework outside the EU [32].

Lineage, hashing and versioning

Lineage tracking via W3C PROV [12] is foundational for any organisation managing provenance at scale, particularly where datasets underpin automated decision-making or model training.

Cryptographic hashing and RFC 3161 trusted timestamping [13] can support governance and due diligence by producing auditable, tamper-evident records, though they attest to data integrity from the point of hashing, not to the accuracy of the underlying measurement.

For B2B licensing and resale, decentralized identifiers (DIDs) and verifiable credentials [14] can support verifiable identity, claims and chain-of-custody evidence, but do not by themselves prove legal ownership or a complete chain of custody.

These controls map onto the hierarchy of trust from Chapter 1: calibration metadata secures Level 1, coordinate standards Level 2, classification metadata Level 3, and lineage and audit frameworks Level 4.

WHICH STANDARD DOES WHICH JOB — AND WHAT NONE OF THEM DOES

FUNCTION	ISO 19115-1 Established	ISO 19157-1 Established	DCAT / GeoDCAT-AP Established	STAC Community	OGC API: Records Established	W3C PROV Established	RFC 3161 Established	DIDs + VCs Established
Describe	✓	•	✓	•	•	•	•	•
Assess quality	•	✓	•	•	•	•	•	•
Catalogue and discover	•	•	✓	✓	✓	•	•	•
Trace lineage	•	•	•	•	•	✓	•	•
Secure and timestamp	•	•	•	•	•	•	✓	•
Assert identity and chain of custody	•	•	•	•	•	•	•	✓
Establish accuracy of the original measurement	No standard covers this							
Establish fitness for the decision	No standard covers this							

Figure 5. The authenticity standards landscape: core standards. See the Appendix for supporting implementation standards (Schema.org, RFC 3161, DIDs and verifiable credentials, ISO/IEC 22989).

Contracts and human review

Standards close the technical gap; contracts and review close the accountability gap. As AI capabilities advance, organisations must improve their governance, validation and accountability practices at a comparable pace. Liability should be defined contractually across the chain (provider, model developer, platform, and end user) rather than left to be disclaimed downstream.

And because non-specialist users increasingly consume geospatial outputs without training in coordinate systems, scale or uncertainty, review processes need a named, competent reviewer before AI-derived data reaches a safety-critical or high-consequence decision.

These expectations are increasingly appearing in procurement, contracts and risk-management practice rather than ather than remaining merely voluntary best practice. In the United States, current GSA guidance for AI acquisition calls for contracts to address issues such as data ownership and intellectual-property rights, access to necessary system components, and vendor handling of government data [15].

Insurance markets are also beginning to examine AI-specific governance and controls in underwriting and coverage design. Liability for AI-generated outputs remains highly context-specific and continues to develop, making contractual clarity an important complement to technical controls.

Flood-map checkpoint: secured

Hashed at creation and timestamped, carrying machine-readable lineage from sensor to classification. A buyer can verify in seconds what produced it and when, though the hash attests to integrity from that point forward, not to the accuracy of the original measurement.

Recommendations

Data producers
Select metadata standards by function: ISO 19115-1 [9] and ISO 19157-1 [16] for metadata and quality; DCAT/GeoDCAT-AP for catalogue interoperability; STAC [17] and OGC API: Records [18] for discovery and access; W3C PROV for lineage.

Outcome: *Matches the standard to the job instead of adopting one standard for everything.*

Platforms
Make W3C PROV lineage, cryptographic hashing, and RFC 3161 timestamping standard practice for data feeding automated decisions.

Outcome: *Produces tamper-evident, auditable records by default.*

Data licensors
Use DIDs and verifiable credentials for B2B licensing where provable chain of custody is a differentiator.

Outcome: *Enables verification without a centralised registry or intermediary.*

Legal and procurement
Require a named, competent human reviewer before AI-derived geospatial data reaches a safety-critical decision.

Outcome: *Closes the gap that pure technical standards cannot: judgment and accountability.*

■ CHAPTER 4

Sovereignty, Transparency and Control

In one sentence:

Sovereignty, licensing and user-facing transparency are not obstacles to trust; handled well, they are how trust crosses borders.

What sovereignty means in practice

Data sovereignty sits at the crossroads of national regulation, organisational responsibility, and the practical realities of modern data ecosystems. It concerns the legal and practical authority governing data, which may depend on far more than where data is collected or stored, a reality that grows more complex as organisations adopt cloud services and operate across borders. Precision matters: data sovereignty (whose law governs), data residency (where data is physically stored), data localisation (requirements to keep it in-country), jurisdiction (which authorities can compel access), and organisational control (who administers it) are distinct concepts. Indigenous and community data sovereignty adds a further dimension increasingly recognised in global geospatial practice.

> **Scope note: digital sovereignty**
> Digital sovereignty is broader than data sovereignty alone. It can extend to the technology stack, supply chains, infrastructure and the corporate entities involved. This report focuses specifically on data sovereignty; the wider digital-sovereignty question may warrant separate treatment.

Satellite imagery adds a distinct sovereignty dimension of its own: collection may occur over one jurisdiction, processing and hosting in another, and distribution or use in several others. Applicable controls may include remote-sensing licences, restrictions on commercial distribution of high-resolution imagery, sensitive-site or resolution controls, national-security rules, and contractual limits on downstream distribution.

Federation as the enabling architecture

Overly strict localisation rules can fragment data ecosystems: if every country stores and describes its datasets differently, cross-border environmental monitoring becomes impossible. Standards like DCAT-AP and GeoDCAT-AP exist precisely to let datasets remain sovereign while staying discoverable and usable across jurisdictions. Federated data infrastructures take this further: data stays within its jurisdiction, but authorised users can run queries or models across distributed datasets, an approach already used in health research where patient data cannot leave national boundaries. Federation is an enabling architecture, not a complete legal solution: queries, embeddings, model parameters and derived outputs may themselves reveal or transfer restricted information and require the same governance attention as the underlying data. UN-GGIM's Integrated Geospatial Information Framework (IGIF) [19] offers a complementary, non-technical model for the same problem: nations retain authority over how their geospatial data is collected, processed and governed, while structured interoperability (common fundamental data themes, shared metadata standards, and the Global Statistical Geospatial Framework [35]) acts as a translation layer, so that nationally sovereign data can still be aligned and compared without forcing every country onto identical systems.

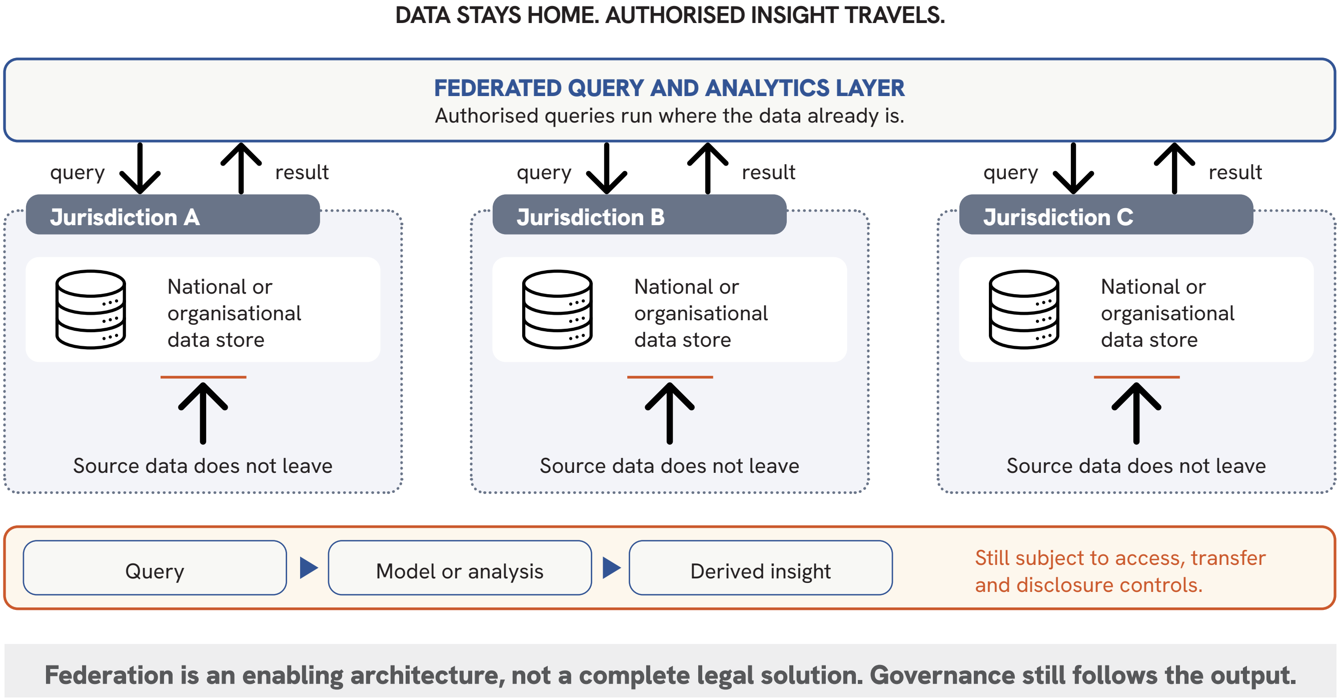

Figure 6. Federated sovereignty: data stays home, insight travels, and queries, embeddings and outputs still carry restricted information.

Flood-map checkpoint: governed

The source pixels remain in-country, while authorised federated processing supports cross-border insurance and planning analysis. Derived outputs and intermediate artefacts remain subject to the same access, transfer and disclosure controls.

From control to disclosure: sovereignty vs. transparency

Sovereignty governs who may control, process and use geospatial data; transparency governs what people must be told when AI transforms, generates or presents it. Together, they address different dimensions of trust in cross-border, AI-enabled geospatial systems.

Transparency of AI-generated geospatial outputs

Since 2 August 2026, Article 50 of the EU AI Act [31] has imposed transparency obligations on certain interactive and generative AI systems, relevant beyond systems classified as high-risk. Providers of AI systems that interact directly with natural persons may need to ensure users are informed of the AI interaction, while providers of systems generating or manipulating synthetic image, audio, video or text content may need to embed effective, machine-readable markings.

Deployers may separately be required to disclose deepfake content or label certain AI-generated text published on matters of public interest. Whether an obligation applies depends on the system, output, intended audience, provider or deployer role, and any applicable exception. The European Commission's Guidelines and FAQ on Article 50 [20][21] are non-binding and provide practical guidance on the Commission's interpretation; authoritative interpretation of the AI Act ultimately rests with the Court of Justice of the European Union. Article 50 applies from 2 August 2026.

A limited grace period applies to AI systems placed on the market before that date, and only in respect of the Article 50(2) marking and detection obligations for AI-generated or manipulated content; providers of those systems must comply with those obligations from 2 December 2026 [21][31], which is directly relevant to GeoAI providers assessing their compliance timeline. Two examples illustrate how these obligations may apply in practice, though coverage always depends on the specific system, output, and role (provider or deployer) involved, and on any applicable exception; it should not be assumed to apply, or assumed not to apply, without that assessment. Digital twins illustrate why provenance and user-facing transparency must operate together. The European Commission's Article 50 Guidelines expressly recognise that AI-generated or manipulated digital twins may fall within the machine-readable marking obligations where their outputs take image, video, audio or text form, subject to applicable industrial, B2B and real-time-content exceptions.

SAR-derived synthetic imagery may raise the same transparency questions as other AI-generated geospatial imagery where the resulting output falls within Article 50's scope. The fact that the source data is SAR does not itself trigger the obligation; what matters is how AI has generated or manipulated the output, how that output is presented or used, and whether the relevant provider or deployer conditions and exceptions apply. Governance-ready systems in both cases should retain source and model lineage internally while enabling appropriate marking or disclosure of synthetic outputs presented to natural persons.

In the United States, AI governance remains distributed across voluntary federal risk-management guidance, agency and sector-specific requirements, procurement controls, and state-level privacy or automated-decision rules. The NIST AI Risk Management Framework [22] is a voluntary cross-sector framework and is currently undergoing revision. For geospatial organisations, applicable provenance and documentation requirements may therefore depend on the use case, customer, sector, contract and jurisdiction, rather than a single horizontal AI statute.

Recommendations

Platforms and governments
Adopt federated data infrastructures, so data remains in-jurisdiction while authorised queries cross borders.

Outcome: *Preserves sovereignty without fragmenting cross-border insight.*

Data producers
Describe datasets with GeoDCAT-AP so they stay sovereign yet discoverable and usable across jurisdictions.

Outcome: *Reduces fragmentation caused by inconsistent national description.*

Policy makers
Back sovereignty with transparent access rules and auditable data-sharing agreements, not blanket localisation.

Outcome: *Keeps data ecosystems usable across borders and sectors.*

Providers, deployers and platforms
Assess GeoAI products against the distinct Article 50 obligations for interactive systems, synthetic-content marking, deepfake disclosure and public-interest text; preserve machine-readable markings through the distribution chain.

Outcome: *Supports clear, timely disclosure at the first interaction or exposure.*

CHAPTER 5

Governance by Design

In one sentence:

Geospatial Foundation Models (GFMs) must ingest, enforce and audit trust: governance built into the architecture, not layered on top.

On 23 July 1972, NASA launched the Earth Resources Technology Satellite, ERTS-1, later renamed Landsat 1. The Landsat programme subsequently became a partnership between NASA and the US Geological Survey, and its global impact expanded dramatically after the USGS adopted a free-and-open-data policy in 2008 [23] [24]. For decades, satellites were largely a government domain; the rise of the commercial satellite industry accelerated the transformation Landsat started. The value is quantifiable: USGS economic studies estimated the annual benefit of Landsat imagery at USD 2.19 billion in 2011 and USD 3.45 billion in 2017 [25], rising to USD 25.6 billion in 2023 [26]. This matters for what follows: the same governance discipline that made archives like Landsat trustworthy and valuable now needs to extend to the models trained on them – otherwise the trust built at the data layer doesn't carry through to the model layer.

Long-running, consistently observed and increasingly open Earth-observation archives such as Landsat are also part of what makes today's GFMs possible. Instead of learning from isolated images, GFMs can be pretrained across large temporal and geographic archives; Prithvi, for example, is pretrained on NASA's Harmonized Landsat and Sentinel-2 dataset. In this sense, GFMs are not a break from the Earth-observation data infrastructure built over the previous five decades, but a new way of extracting and compressing knowledge from it.

GFMs are where the shift from measurement-based to model-based trust, described throughout this report, becomes concrete. Earlier chapters addressed provenance, standards, and governance for individual datasets; a GFM compresses many such datasets into a single trained model whose internal representations and transformations are not directly inspectable in the same way as the source datasets from which it was trained. The trust questions raised so far in this report, about origin, integrity, uncertainty, and accountability, do not disappear at this stage. They become harder to observe, and harder to correct after the fact. Understanding what GFMs do, and what governing them requires, is therefore central to this report's argument: provenance and metadata discipline earlier in the pipeline will not be sufficient if the model layer itself cannot be governed.

GFMs ingest satellite and aerial imagery, LiDAR, SAR, DEMs, vectors, textual metadata, and mobility data. This heterogeneity makes governance harder and more essential: labels are sparse, ground truth is limited, and much of the training is unsupervised at a global scale. Several GFMs illustrate the range of approaches: SkySense combines multi-modal, multi-resolution and multi-temporal inputs and reports state-of-the-art results across Earth-observation benchmarks [27]; Prithvi, developed by NASA and IBM, is pretrained on NASA's Harmonized Landsat and Sentinel-2 data and released as an open model [28]; and the Clay Foundation Model is an independently governed open-source GFM trained across multiple public satellite and aerial sources [29]. Strong benchmark performance is evidence of technical capability, not by itself proof of suitability for regulatory or decision-grade deployment. If such models mature as expected, they may become the default substrate for geo-analytics and automated geo-reasoning; governance must be designed in from the outset.

What to ask before trusting a GFM

GFM developers should document how their training data was selected, curated and transformed. Such documentation is a starting point for establishing trust, not a substitute for it. Before relying on a GFM output, users should be able to answer three questions:

- What data was the model trained on, and how was that data manipulated?
- What temporal period is represented in the training data, and where is the model being asked to interpolate or extrapolate beyond it?
- How does the model perform against my specific benchmarks, not just published leaderboards?

It is worth being explicit about what a GFM does to geospatial data: it compresses spatial context into a far more succinct representation. That representation is model-dependent rather than neutral: it reflects choices about training data, preprocessing, model architecture, objectives and evaluation. Reliance on the resulting representation does not by itself transfer or eliminate responsibility for decisions built on top of it.

Caution: model collapse

When recursively generated synthetic data contaminates or progressively replaces the original data distribution used to train the next generation of models, it can contribute to model collapse and the degradation of provenance across successive training cycles [30].

Flood-map checkpoint: ingested

Before the flood-map data is used to train a GFM, the authenticity and sovereignty gates check its lineage, licensing and jurisdiction. Its source and model lineage are recorded in the governance pipeline, allowing later outputs to be linked to the relevant model version, input-data classes and applicable controls.

A governance-ready GFM architecture

Metadata is not merely stored: it is queried by GFMs, validated by governance engines, and audited by regulators. The architecture below organises this into three layers, with a single audit capability spanning all of them rather than appearing only at the end of the pipeline.

This point matters more than it may first appear

Traditional geospatial accuracy assessments generally quantify observed error against reference data using defined statistical measures and test procedures. A GFM's confidence or uncertainty measure instead characterises uncertainty in a model prediction and may itself require calibration and validation. The two therefore answer different questions and should not be treated as interchangeable measures of trust.

GFM governance must evaluate both predictive performance against appropriate reference data and the calibration of reported uncertainty. GFMs need appropriate uncertainty-quantification methods, including stochastic or Bayesian approaches where suitable, to support uncertainty awareness and robustness to missing or noisy data: geospatial data is incomplete, sensors disagree, labels are sparse, and decisions have consequences.

GOVERNANCE IS IN THE ARCHITECTURE, NOT ON TOP OF IT

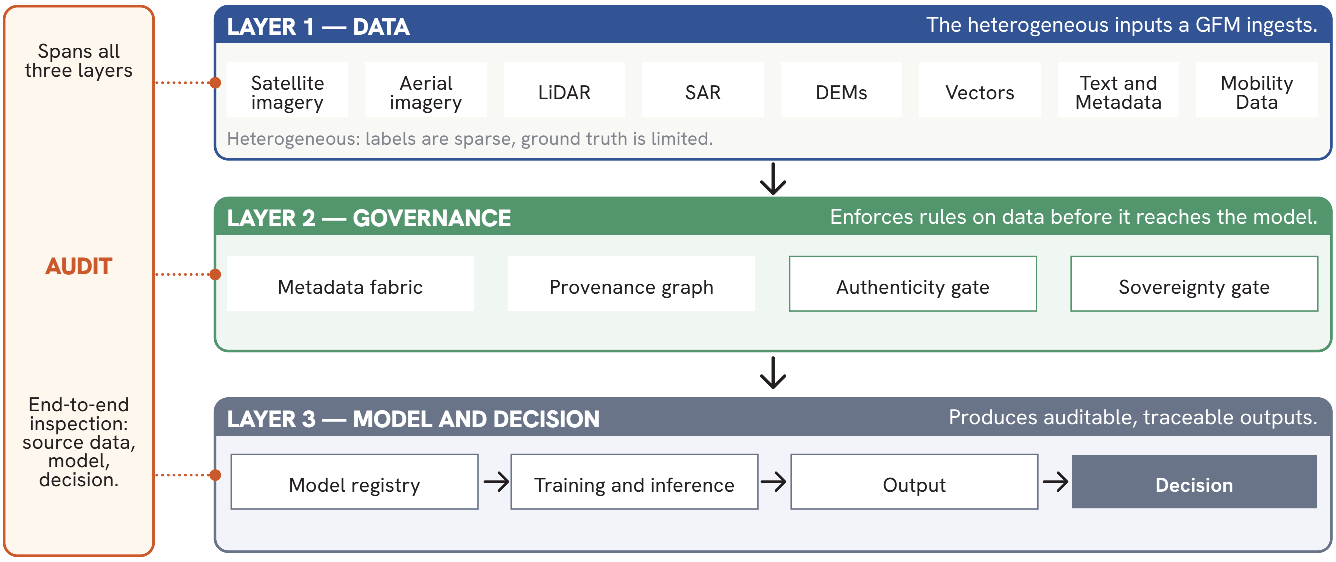

Audit is not performed on the output. It runs through every layer.

Figure 7. The governance-ready GFM architecture: three layers, with audit spanning all of them.

Recommendations

GFM developers
Treat metadata as something GFMs query, validate, and audit, not something merely stored: sensor physics, acquisition context, transformation lineage, and model-specific metadata.

Outcome: *Makes metadata an active part of the model, not passive documentation.*

GFM developers
Support GFMs with stochastic and Bayesian methods for uncertainty awareness and robustness to sparse, noisy, or missing data.

Outcome: *Gives models a principled way to express what they do not know.*

Legal and policy
Apply existing geospatial law early to GFM development and deployment, including remote-sensing regulation, geodata licensing, location privacy, national mapping restrictions and responsibility for spatially consequential outputs.

Outcome: *Avoids treating GFM governance as legally novel where established geospatial rules already apply.*

CHAPTER 6

WGIC Action Agenda

In one sentence:

Trust will be rebuilt through provenance, governance, and accountability at industry scale, and no single stakeholder can do it alone.

This report's core message is that the geospatial sector is moving from measurement-based trust to model-based trust, and that trust has to be rebuilt on the four foundations set out in the Executive Summary: provenance, governance, metadata, and accountability. Chapters 1–5 set out why, and what that looks like in the data lifecycle, sovereignty, and GFM architecture. This chapter turns that into an agenda: who acts, in what order, and what WGIC should lead.

A three-horizon roadmap

The single governance umbrella referenced above is not a new idea introduced here: it is the governance layer from Chapter 5's GFM architecture, applied at industry scale: a shared metadata fabric, a provenance graph, a policy engine, a model registry, and an audit layer that lets regulators inspect provenance end to end.

THE WHOLE AGENDA, IN ABOUT TWENTY SECONDS

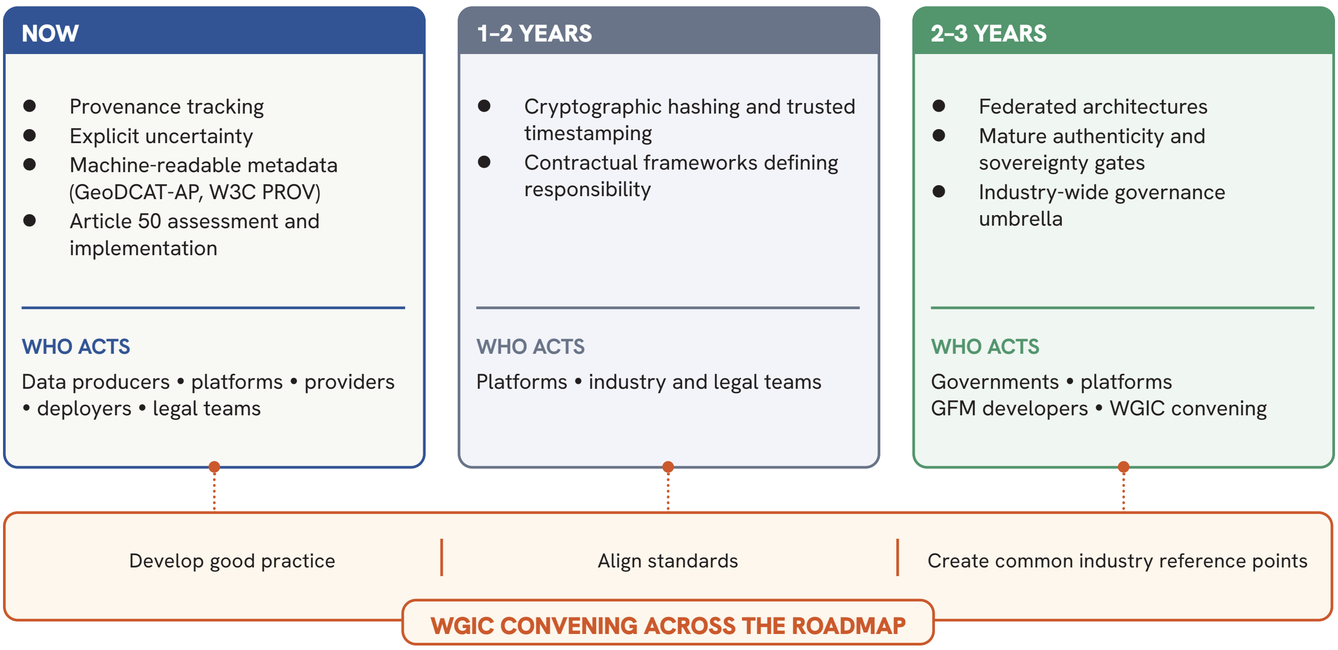

Figure 8. The WGIC action agenda: who acts, and when.

This governance umbrella supports transparency and enforceable standards while reducing duplication and inconsistent preprocessing, and improving the detection, traceability and management of erroneous or hallucinated outputs. None of this removes the complexity of fast-changing legislation, or the legal risk attached to specific use cases.

The business case

Provenance and governance should not be read only as compliance costs. In the GeoAI era they are strategic enablers of trust, adoption and value creation: organisations that can demonstrate a dataset's origin, integrity and lineage can differentiate themselves commercially, may be better positioned for partnerships and procurement processes that require verifiable provenance, and can strengthen customer confidence.

There is an obvious objection: provenance infrastructure costs money and adds friction, while buyers do not always pay an identifiable premium for verified data. The business case will vary by sector and consequence level. Organisations should weigh implementation cost against regulatory exposure, contractual liability, rework, reputational harm, and the cost of decisions that cannot later be reconstructed or defended. Building provenance controls during system design is generally more efficient than retrofitting them after a failure or dispute.

Maturity checklist

Where does your organisation stand?

1. Can you trace any dataset you rely on back to its original sensor and calibration state?
2. Do you know which version of a dataset informed a given decision?
3. Is your metadata machine-readable and standards-compliant (GeoDCAT-AP, W3C PROV)?
4. Could you prove chain of custody for your licensed data, cryptographically?
5. Do your AI-derived products carry explicit uncertainty information?
6. Do your contracts define who is liable when derived data is wrong?

Mostly "no"? Chapter 3 is your starting point.

The path forward demands industry-wide coordination. No single stakeholder can ensure trust alone. WGIC, positioned across every stage of the global geospatial data value chain, can lead the development of best practices, convene expertise, and champion shared standards. Trust in geospatial data can be rebuilt, not by returning to past methods, but by focusing on well-defined use cases underpinned by modern systems that deliver transparency, accountability, and interpretability at scale.

Concretely, this depends on three distinct but complementary roles. Industry, organised through bodies like WGIC, must build transparent, auditable pipelines and treat provenance and metadata as core infrastructure, not an afterthought. Governments must set the governance frameworks, standards, and liability rules that give model-based outputs legal and institutional weight, especially where decisions affect public safety or policy. And academia must provide the independent scrutiny, methodological rigour, and training that validate AI-driven methods and prepare the next generation of practitioners. Trust in geospatial data is rebuilt through this three-way collaboration (industry building it, government regulating it, academia validating it), not by any one actor alone.

It is worth restating plainly: none of the standards, gates, or architectures in this report create trust by themselves. They create the conditions under which people (reviewers, data producers, decision-makers) can extend trust responsibly. Trust is ultimately a human judgment, informed by evidence; it is not a property that machines or standards possess on their own.

WGIC's role

WGIC convenes members across every stage of the geospatial value chain (data producers, platforms, GFM developers, and policy teams) to turn this agenda into shared practice: developing best practices, aligning on standards, and giving members a common reference point as regulation and AI capability both move quickly.

Supporting implementation standards

These standards support the core standards shown in Figure 5 and referenced in Chapter 3; they have been moved here to keep the main chapter focused.

Table 2. Supporting implementation standards.

Standard	What it does	Maturity
Schema.org	Can improve the discoverability of geospatial assets within broader data ecosystems; does not itself certify quality or licensing.	Established
ISO/IEC 22989	Establishes common concepts and terminology for AI; provides a shared vocabulary rather than a geospatial metadata-readiness framework.	Published
RFC 3161	Trusted timestamping; produces tamper-evident records from the point of hashing.	Established
DIDs + verifiable credentials	Can support verifiable identity, claims, and chain-of-custody evidence; do not by themselves prove legal ownership.	Established W3C standards; geospatial adoption developing
ISO 19119	Geographic information services: describes how geospatial services (not just data) are catalogued and discovered.	Established (ISO/TC 211)
ISO 19110	Methodology for feature cataloguing: defines feature types, attributes and relationships within a dataset.	Established (ISO/TC 211)
ISO/IEC 25024	Data quality measurement: defines measures for evaluating data quality characteristics.	Established (ISO/IEC JTC 1)
ISO 5725	Accuracy (trueness and precision) of measurement methods and results.	Established (ISO)
FGDC metadata standards	US federal profile; FGDC endorses ISO 19115-1 and related ISO standards while continuing to support the legacy CSDGM standard.	Established (US FGDC)

Feature- and service-level cataloguing (ISO 19119, ISO 19110) [33] [34] matters more, not less, in an AI-driven pipeline: AI systems need this metadata to assess fitness for use, but populating it well increasingly requires AI assistance itself, since the pool of specialist subject-matter experts available to do it manually is shrinking relative to the volume of data being produced.

References

1 ISO/IEC 25024:2015, Systems and software engineering, SQuaRE: measurement of data quality. https://www.iso.org/standard/35749.html

2 ISO 5725-1:2023, Accuracy (trueness and precision) of measurement methods and results, part 1: general principles and definitions. https://www.iso.org/standard/69418.html

3 Copernicus Emergency Management Service. Global Flood Monitoring (GFM), technical information. European Commission. https://global-flood.emergency.copernicus.eu/technical-information/glofas-gfm/

4 Microsoft. Global ML Building Footprints (living repository; accessed August 2026). https://github.com/microsoft/GlobalMLBuildingFootprints

5 Sirko, W. et al. (2021). Continental-scale building detection from high-resolution satellite imagery. Google Research Open Buildings. https://sites.research.google/gr/open-buildings/

6 W3C. DCAT 3: Data Catalog Vocabulary. https://www.w3.org/TR/vocab-dcat-3/

7 European Commission, SEMIC. GeoDCAT-AP 3.1.0. https://semiceu.github.io/GeoDCAT-AP/releases/3.1.0/

8 Schema.org. Official documentation. https://schema.org/

9 ISO 19115-1:2014, Geographic information, metadata, part 1: fundamentals. ISO/TC 211. https://www.iso.org/standard/53798.html

10 ISO/IEC 22989:2022, Artificial intelligence: concepts and terminology. https://www.iso.org/standard/74296.html

11 Federal Geographic Data Committee. FGDC-STD-001-1998, Content Standard for Digital Geospatial Metadata (CSDGM); FGDC news on ISO metadata standards endorsement. https://www.fgdc.gov/standards/news/fgdc-iso-metadata-standards

12 W3C. PROV-O: The PROV Ontology. https://www.w3.org/TR/prov-o/

13 IETF RFC 3161, Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP). https://datatracker.ietf.org/doc/html/rfc3161

14 W3C (2022). Decentralized Identifiers (DIDs) v1.0. https://www.w3.org/TR/did-1.0/
W3C (2025). Verifiable Credentials Data Model 2.0. https://www.w3.org/TR/vc-data-model-2.0/

15 U.S. General Services Administration (2026). CIO 2185.1C, Accelerating Responsible Use of Artificial Intelligence at GSA (current directive). https://www.gsa.gov/directives-library/accelerating-responsible-use-of-artificial-intelligence-at-gsa

16 ISO 19157-1:2023, Geographic information, data quality, part 1: general requirements. ISO/TC 211. https://www.iso.org/standard/78900.html

17 OGC. SpatioTemporal Asset Catalog (STAC) 1.1, OGC Community Standard 25-004. https://docs.ogc.org/cs/25-004/25-004.html

18 OGC. OGC API - Records - Part 1: Core, OGC Standard 20-004r1 (normative specification). https://docs.ogc.org/is/20-004r1/20-004r1.html

19 United Nations Committee of Experts on Global Geospatial Information Management (UN-GGIM). United Nations Integrated Geospatial Information Framework (UN-IGIF). https://ggim.un.org/UN-IGIF/overview/

20 European Commission (2026). Guidelines on transparency obligations for providers and deployers of certain AI systems under Article 50 of Regulation (EU) 2024/1689, C(2026) 5054 final, 20 July 2026. https://digital-strategy.ec.europa.eu/en/policies/guidelines-transparency-ai-generated-content

21 European Commission. FAQ: Transparency obligations under Article 50 of the AI Act. https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act

22 National Institute of Standards and Technology (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. https://www.nist.gov/itl/ai-risk-management-framework

23 NASA Science. Landsat 1 mission history. https://science.nasa.gov/mission/landsat-1/

24 U.S. Geological Survey (2018). Free, Open Landsat Data Unleashed the Power of Remote Sensing a Decade Ago. https://www.usgs.gov/news/free-open-landsat-data-unleashed-power-remote-sensing-a-decade-ago

25 Straub, C.L., Koontz, S.R., and Loomis, J.B. (2019). Economic valuation of Landsat imagery. U.S. Geological Survey Open-File Report 2019-1112. https://doi.org/10.3133/ofr20191112

26 U.S. Geological Survey (2024). Economic Valuation of Landsat and Landsat Next 2023. https://www.usgs.gov/media/files/economic-valuation-landsat-and-landsat-next-2023

27 Guo, X. et al. (2024). SkySense: A multi-modal remote sensing foundation model towards universal interpretation for Earth observation imagery. CVPR 2024. https://openaccess.thecvf.com/content/CVPR2024/html/Guo_SkySense_A_Multi-Modal_Remote_Sensing_Foundation_Model_Towards_Universal_Interpretation_CVPR_2024_paper.html

28 NASA IMPACT and IBM Research (2023). HLS Geospatial Foundation Model (Prithvi), trained on the Harmonized Landsat and Sentinel-2 dataset. https://www.earthdata.nasa.gov/news/nasa-ibm-openly-release-geospatial-ai-foundation-model-nasa-earth-observation-data

29 Clay Foundation (2024). Clay Foundation Model: an open-source AI model for Earth, v1.5. https://github.com/Clay-foundation/model

30 Shumailov, I. et al. (2024). AI models collapse when trained on recursively generated data. Nature, 631(8022), 755-759. https://www.nature.com/articles/s41586-024-07566-y

31 Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), consolidated text incorporating the Digital Omnibus transitional amendments. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

32 ANZLIC - the Spatial Information Council. AS/NZS ISO 19115.1:2015 Metadata. https://www.anzlic.gov.au/resources/asnzs-iso-1911512015-metadata

33 ISO 19119:2016, Geographic information - Services. ISO/TC 211. https://www.iso.org/standard/59221.html

34 ISO 19110:2016, Geographic information - Methodology for feature cataloguing. ISO/TC 211. https://www.iso.org/standard/57303.html

35 United Nations Committee of Experts on Global Geospatial Information Management (UN-GGIM). The Global Statistical Geospatial Framework (GSGF). https://unstats.un.org/wiki/display/GSGF

Acknowledgements

This report was developed under the WGIC Data Committee, with review, feedback and technical input from the following committee members, and with support from the WGIC secretariat and editorial team.

WGIC Data Committee

John Renard – Chair, Data Committee, Airedale Advisory Limited

James Van Rens – Senior Vice President, RIEGL International

Philip Chisholm – Manager Software Engineering, TomTom

Chantal Jorawsky – Product Marketing Manager, Geospatial, Trimble

James Banting – VP Research and Strategic Projects, Sparkgeo

Deniz Karagulle – Global Community Engagement Technical Lead, Esri

Kathleen Martin – Geospatial Services and Data Analytics Practice Lead, MBS

John Bowers – Head Geointelligence New Business, MDA

Dr. Ramesh Kajrolkar – COO & Chief of International Partnerships, AI-InfraSolutions

Arjen Crince – Technical Director, iSpatial Techno Solutions (IST)

Guillermo Gutiérrez – Geospatial Technical Manager, IIC Technologies Europe

Kevin Pomfret – Partner, Pierson Ferdinand LLP

Rachel Tidmarsh – Chief Executive Officer, Bluesky, Woolpert

Rob Milicevic – Director, Glass Aftercare Ltd

WGIC secretariat

Aaron Addison – Executive Director, WGIC

Helen Gilmartin – Associate Director, WGIC

Bhanu Rekha – Communication Strategist, WGIC

Aarti Iyer – Executive Assistant, WGIC

Editing and consolidation

Kuhelee Chandel – Associate Director – Research, WGIC

World Geospatial Industry Council
Fluwelen Burgwal 58,
2511 CJ The Hague
The Netherlands
Email: info@WGICouncil.org

Visit us: WGICouncil.org