Data Modelling in the Digital Twin of Madrid City

From asset management to observability

David Fernandez (Libelium) | **Juanje Muñoz** (Madrid Digital Office) | **Dr. Antonio Jara** (Libelium)

2nd AIOTI Annual Workshop on Semantic Interoperability for Digital Twins

Behind the change. Beyond the challenge.

The Aspiration: Data-Driven Datocracy

A dynamic, virtual reflection of the city built on absolute data reliability.

- **Data quality** is the foundational requirement for municipal ‘datocracy’.
- The goal is a **living operational tool**, not merely a static 3D map.

The Obstacle: An Identity Crisis in the IoT Stack

Ambiguous IDs prevent reliable macroscopic insights

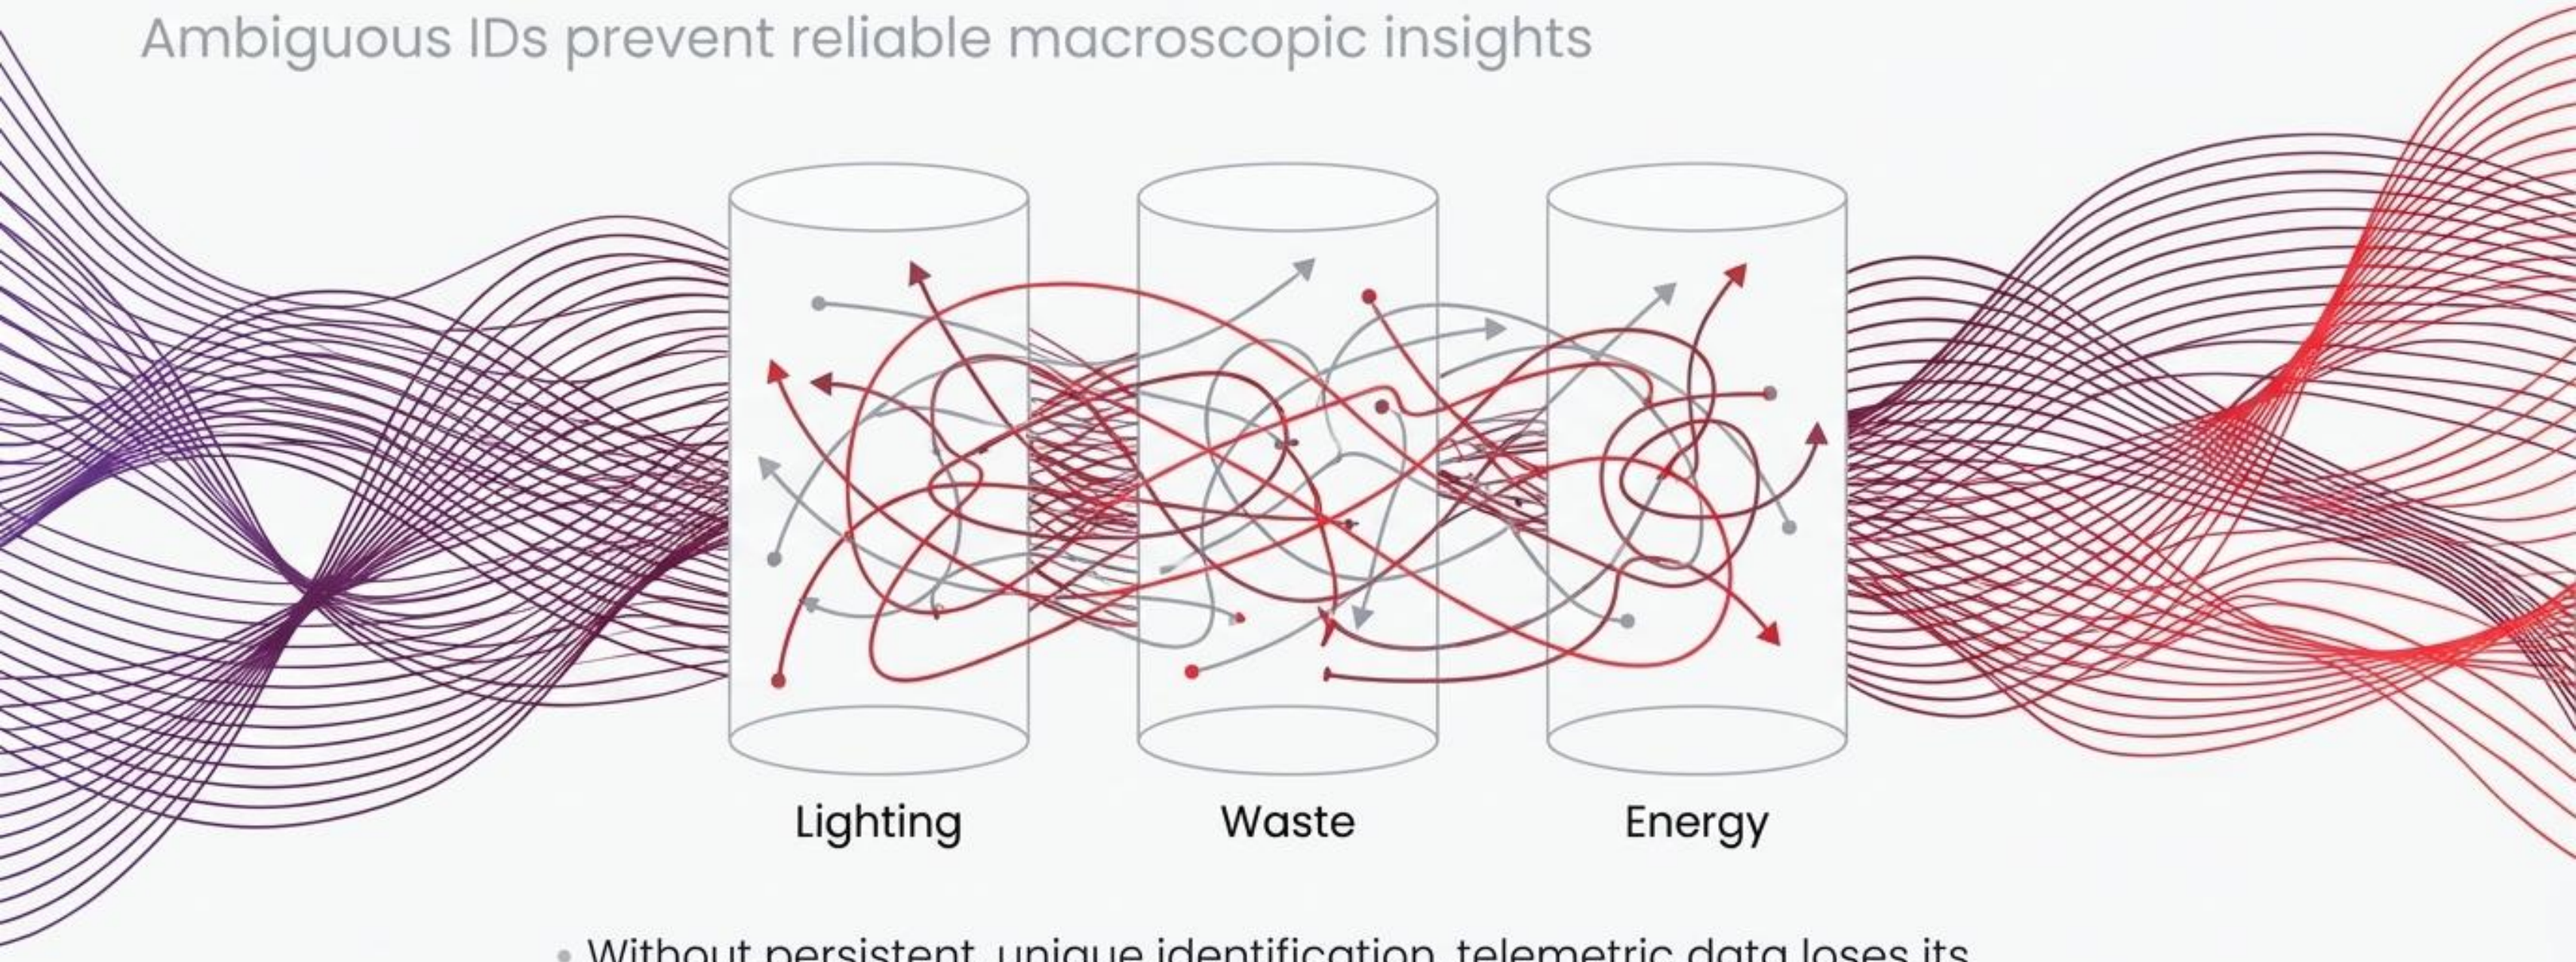

- Without persistent, unique identification, telemetric data loses its physical and historical context over time.

The Obstacle: The Lack of Unique Identity

We cannot unambiguously identify the primary origin of every data point.

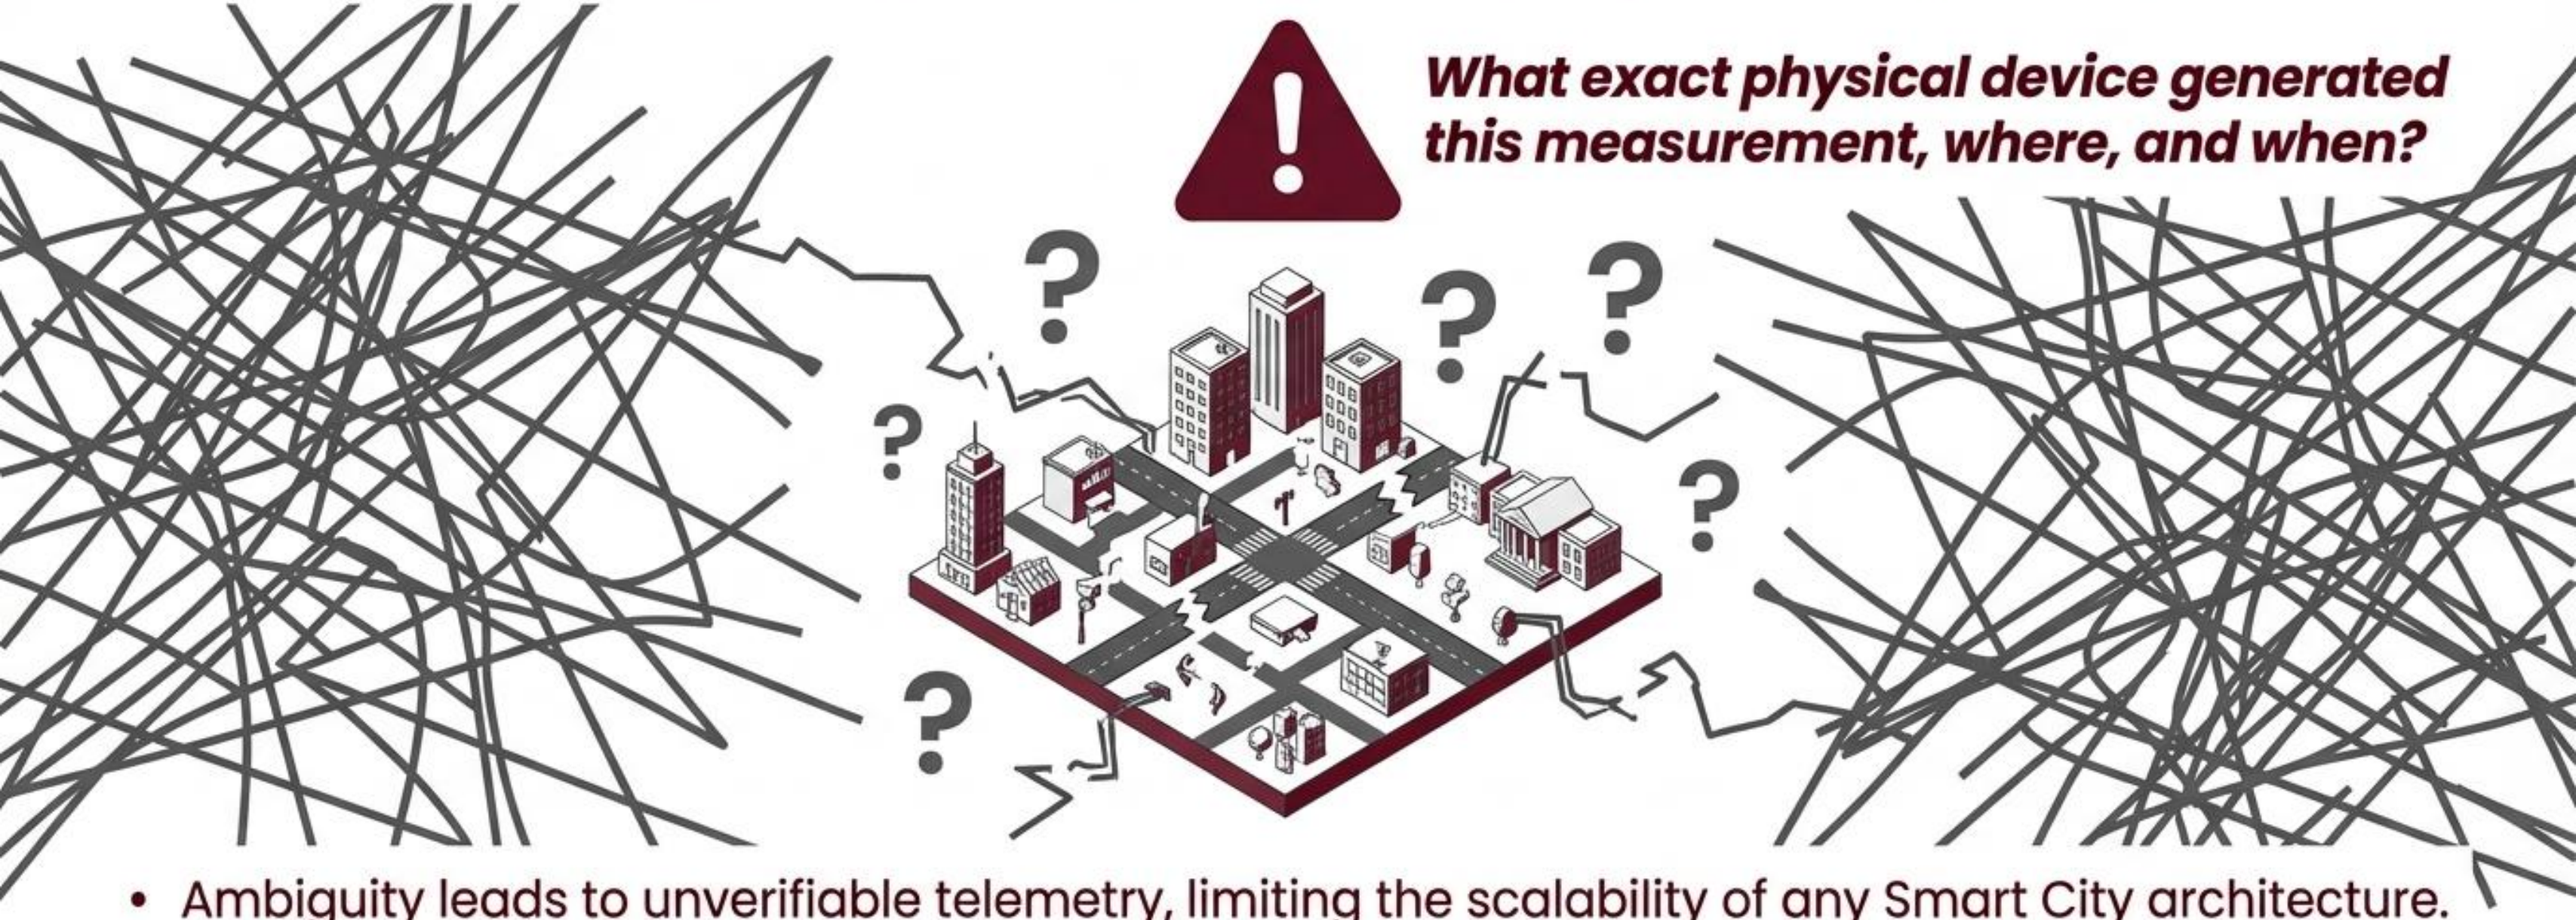

- Ambiguity leads to unverifiable telemetry, limiting the scalability of any Smart City architecture.
- Without strict origin tracing, building a rigorous, trustworthy Digital Twin is impossible.

The Four Pillars of Urban Data Chaos

Ambiguous IDs

Different municipal systems utilize internal identifiers that inevitably collide.

Disconnected Context

Telemetry data arrives stripped of precise physical location, maintenance contracts, or lifecycle status.

Lack of Persistence

Historical data continuity is destroyed when a physical asset is moved or system-managed.

Data Silos

Vertical services operate on proprietary naming schemas, blocking transversal city views.

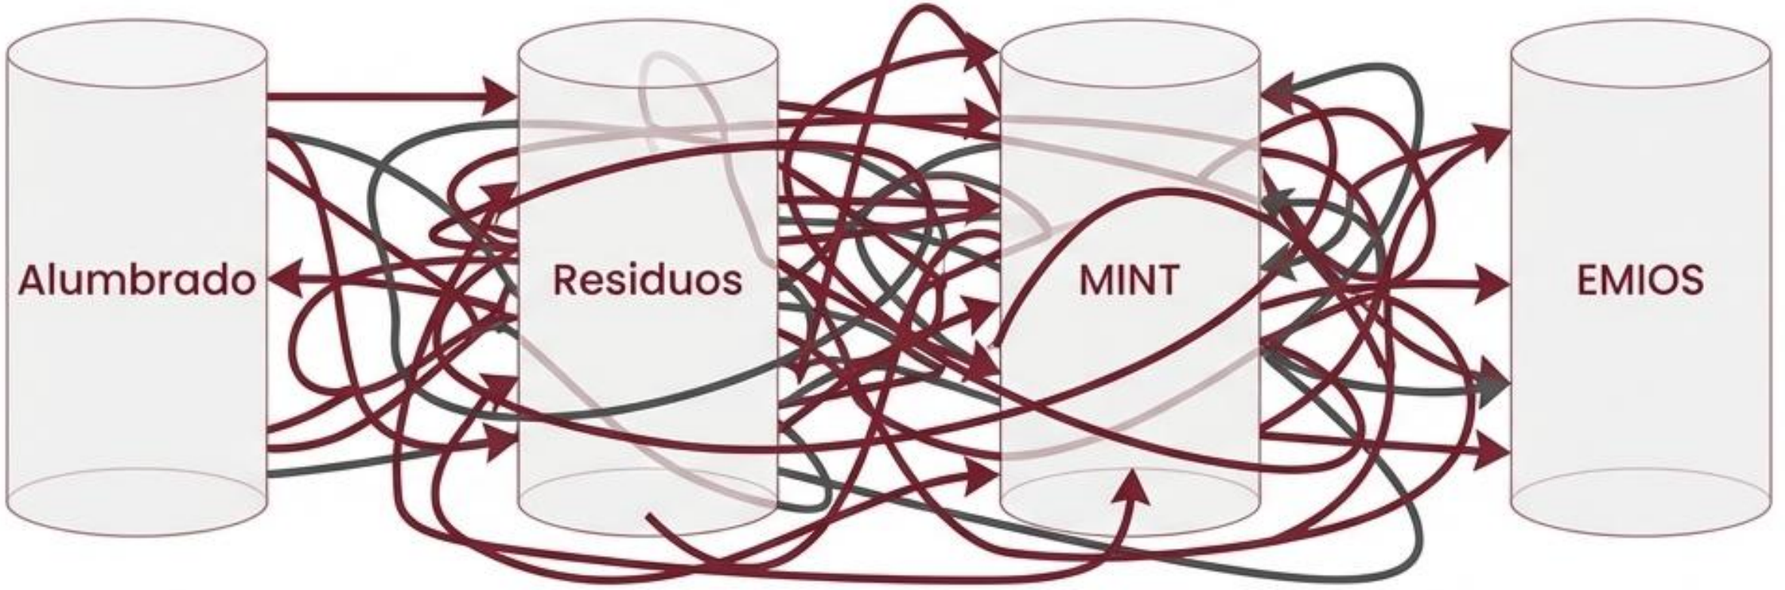

The Disconnect Between Device and Location

Asset identity must survive hardware replacements.

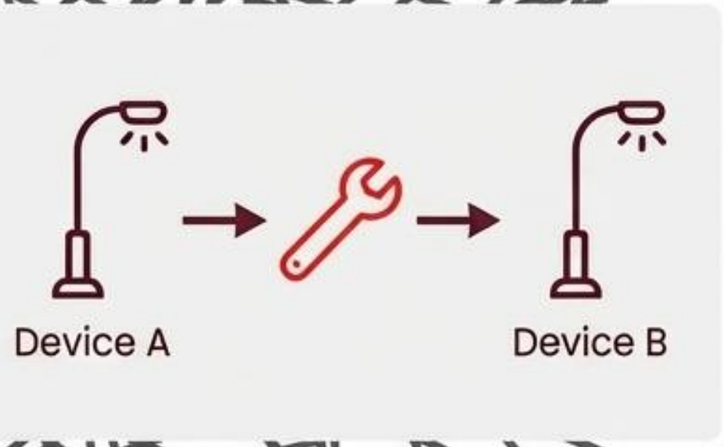

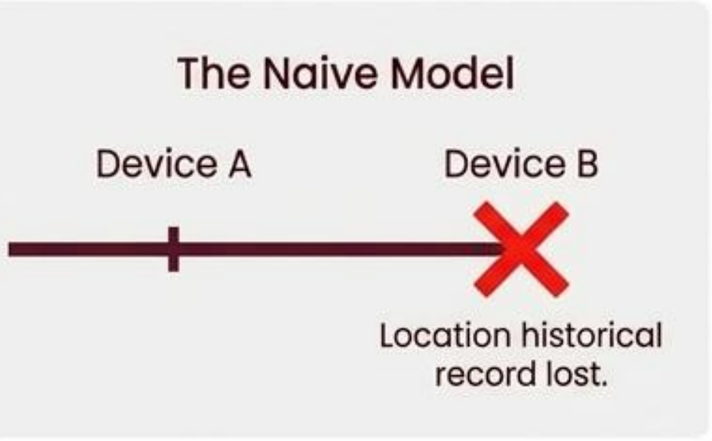

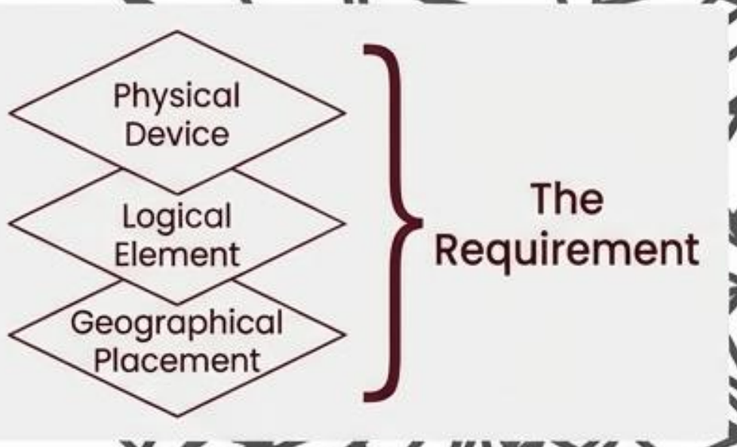

- Every day, thousands of technicians from dozens of companies alter urban elements.
- Data queries must target the permanent physical location, not the transient piece of hardware.

Unfolding Legacy Flat Records

Separating the physical asset from the functional location

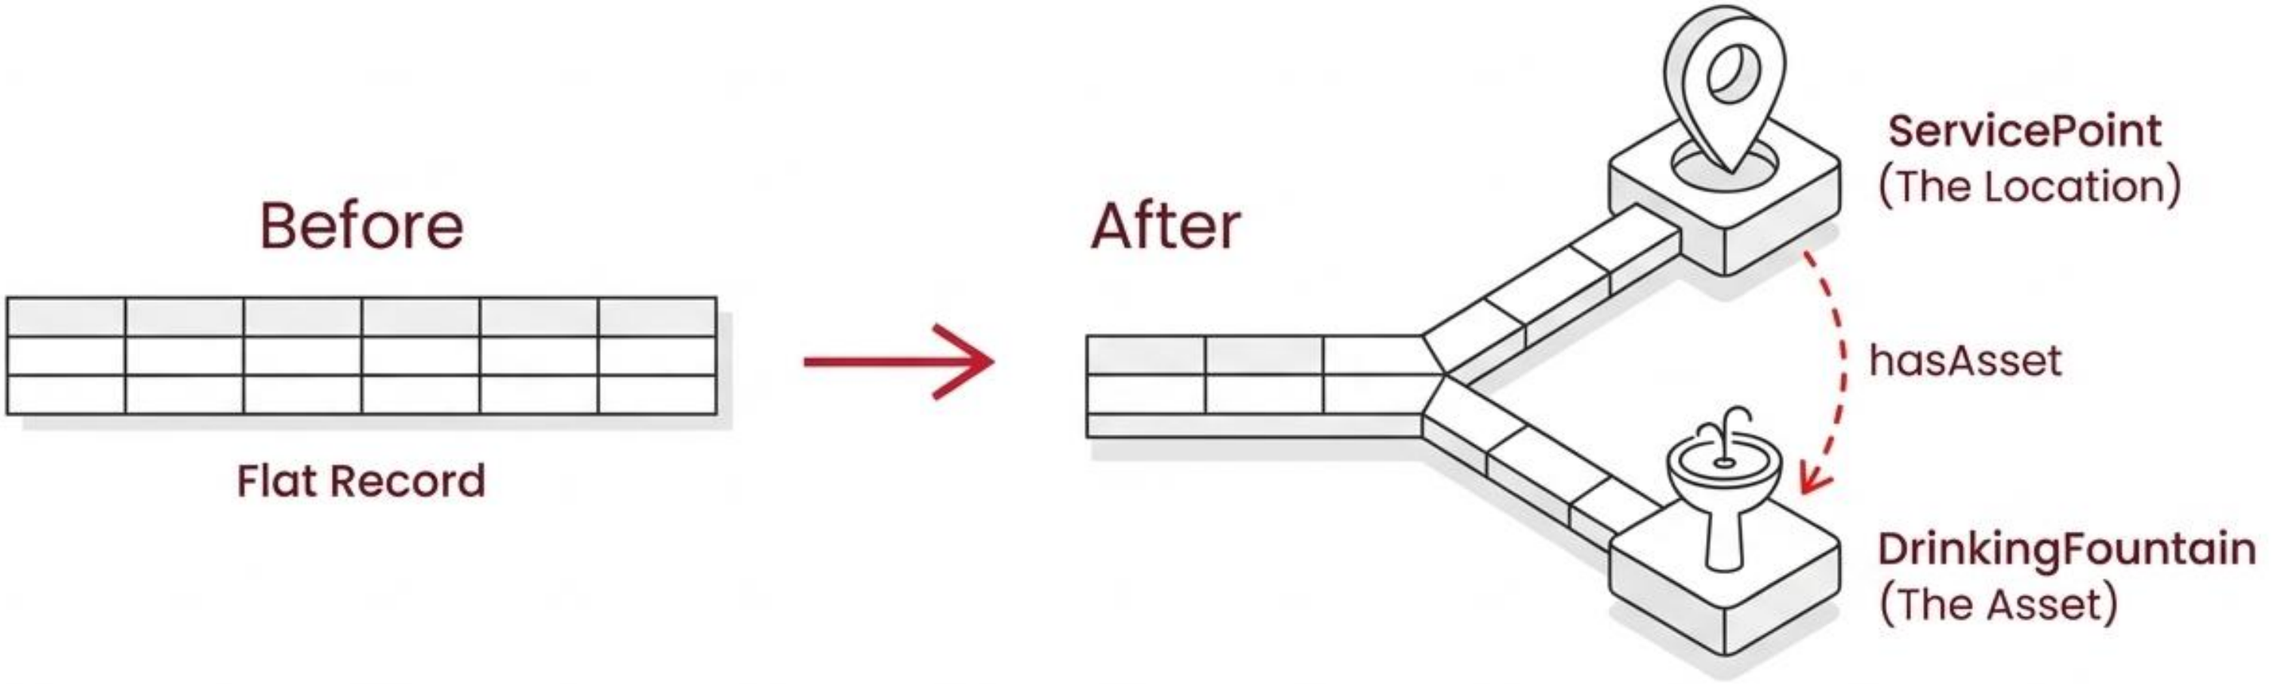

- Legacy systems merge hardware, location, and logic into one ambiguous ID, breaking traceability when hardware changes.

The Paradigm Shift: Building a Knowledge Graph

Context defined by dynamic relationships, not hardcoded strings

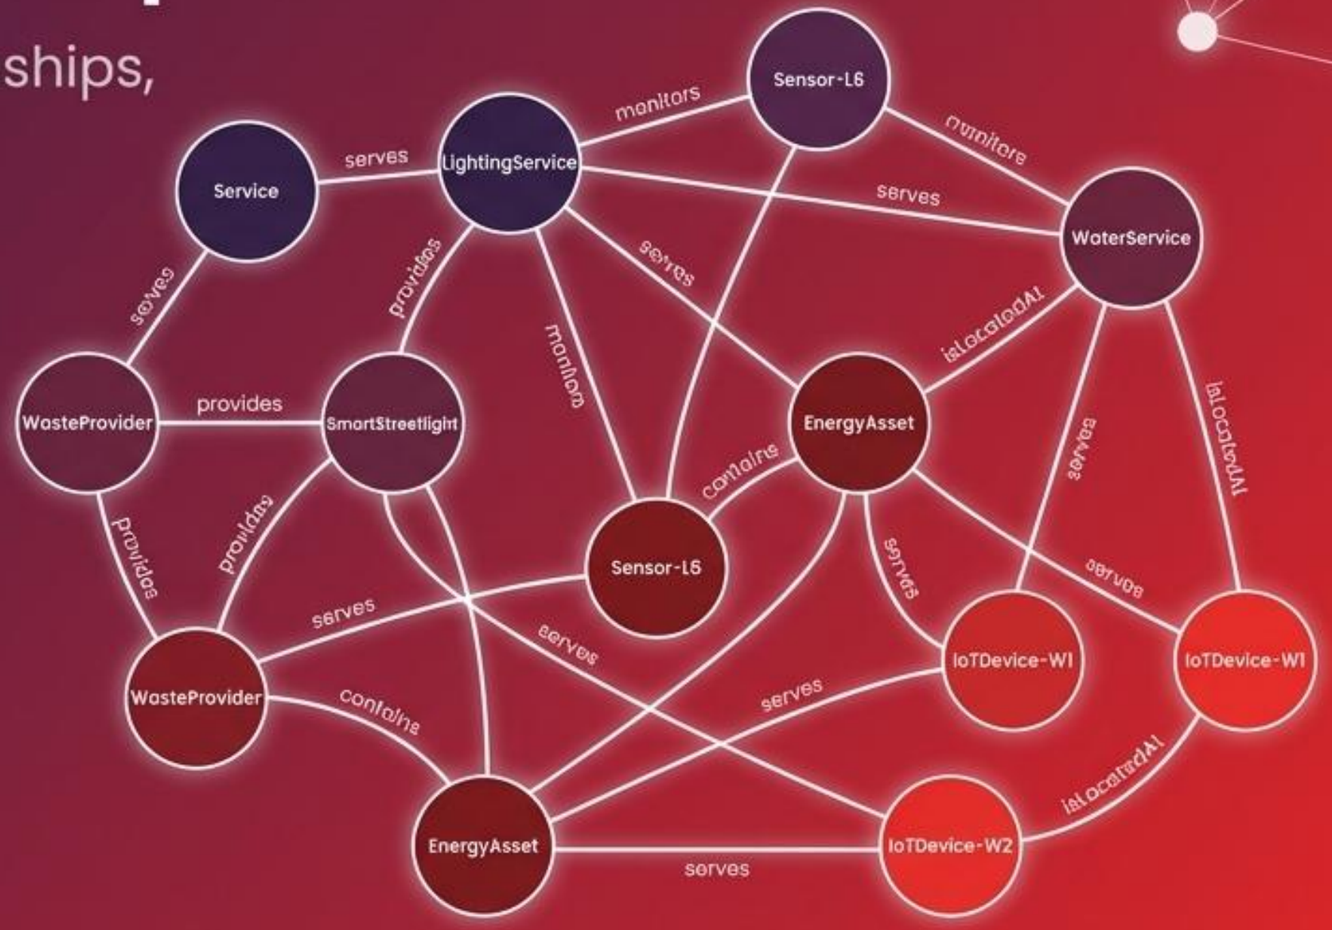

- Hierarchy and context are managed via dynamic links (NGSI-LD), creating a fully navigable ecosystem.

Behind the change. Beyond the challenge.

Intelligent Asset Management

The definitive cornerstone of the Madrid Digital Twin.

- Establishes the indispensable semantic and relational baseline required for a scalable urban operating system.
- Provides a unified vision of the asset lifecycle, optimizing maintenance, operations, and service level agreements (SLAs).

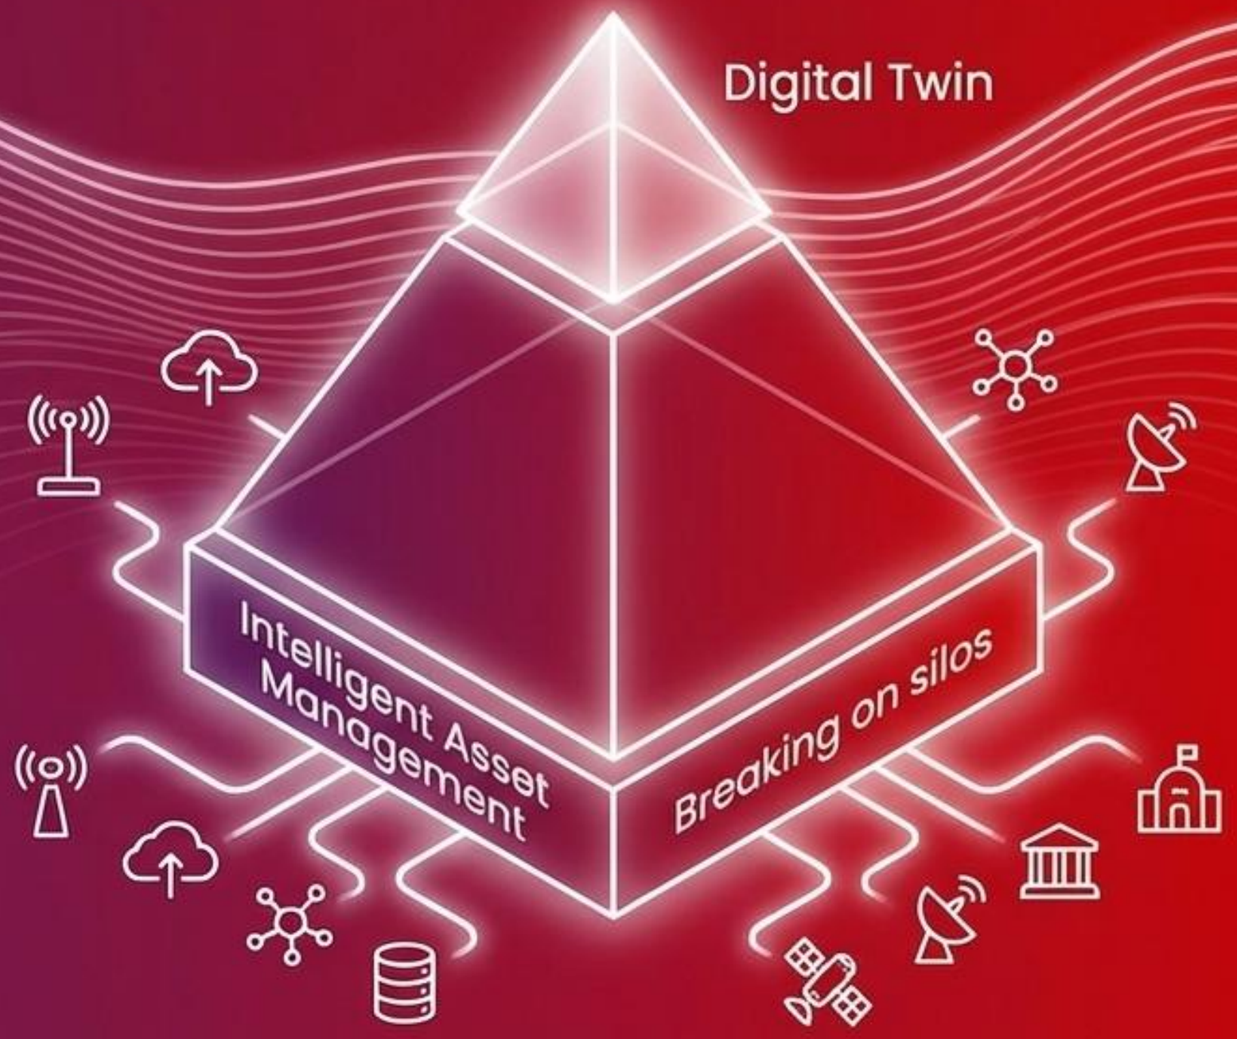

The Architecture of Observability

The Datocracy Equation

Standardized Transport

HTTPS / JSON / XML payload delivery.

Semantic Context

The @context layer.

Persistent Identity

Unambiguous digital tracking.

=

The City Operating System (UrbanOS)

- Breaks vertical data silos by mandating standard integration protocols for all municipal providers.
- Transforms raw, ambiguous telemetry into semantically consistent, actionable knowledge.

@context is all we need

Resolving ambiguity through standard ontologies.

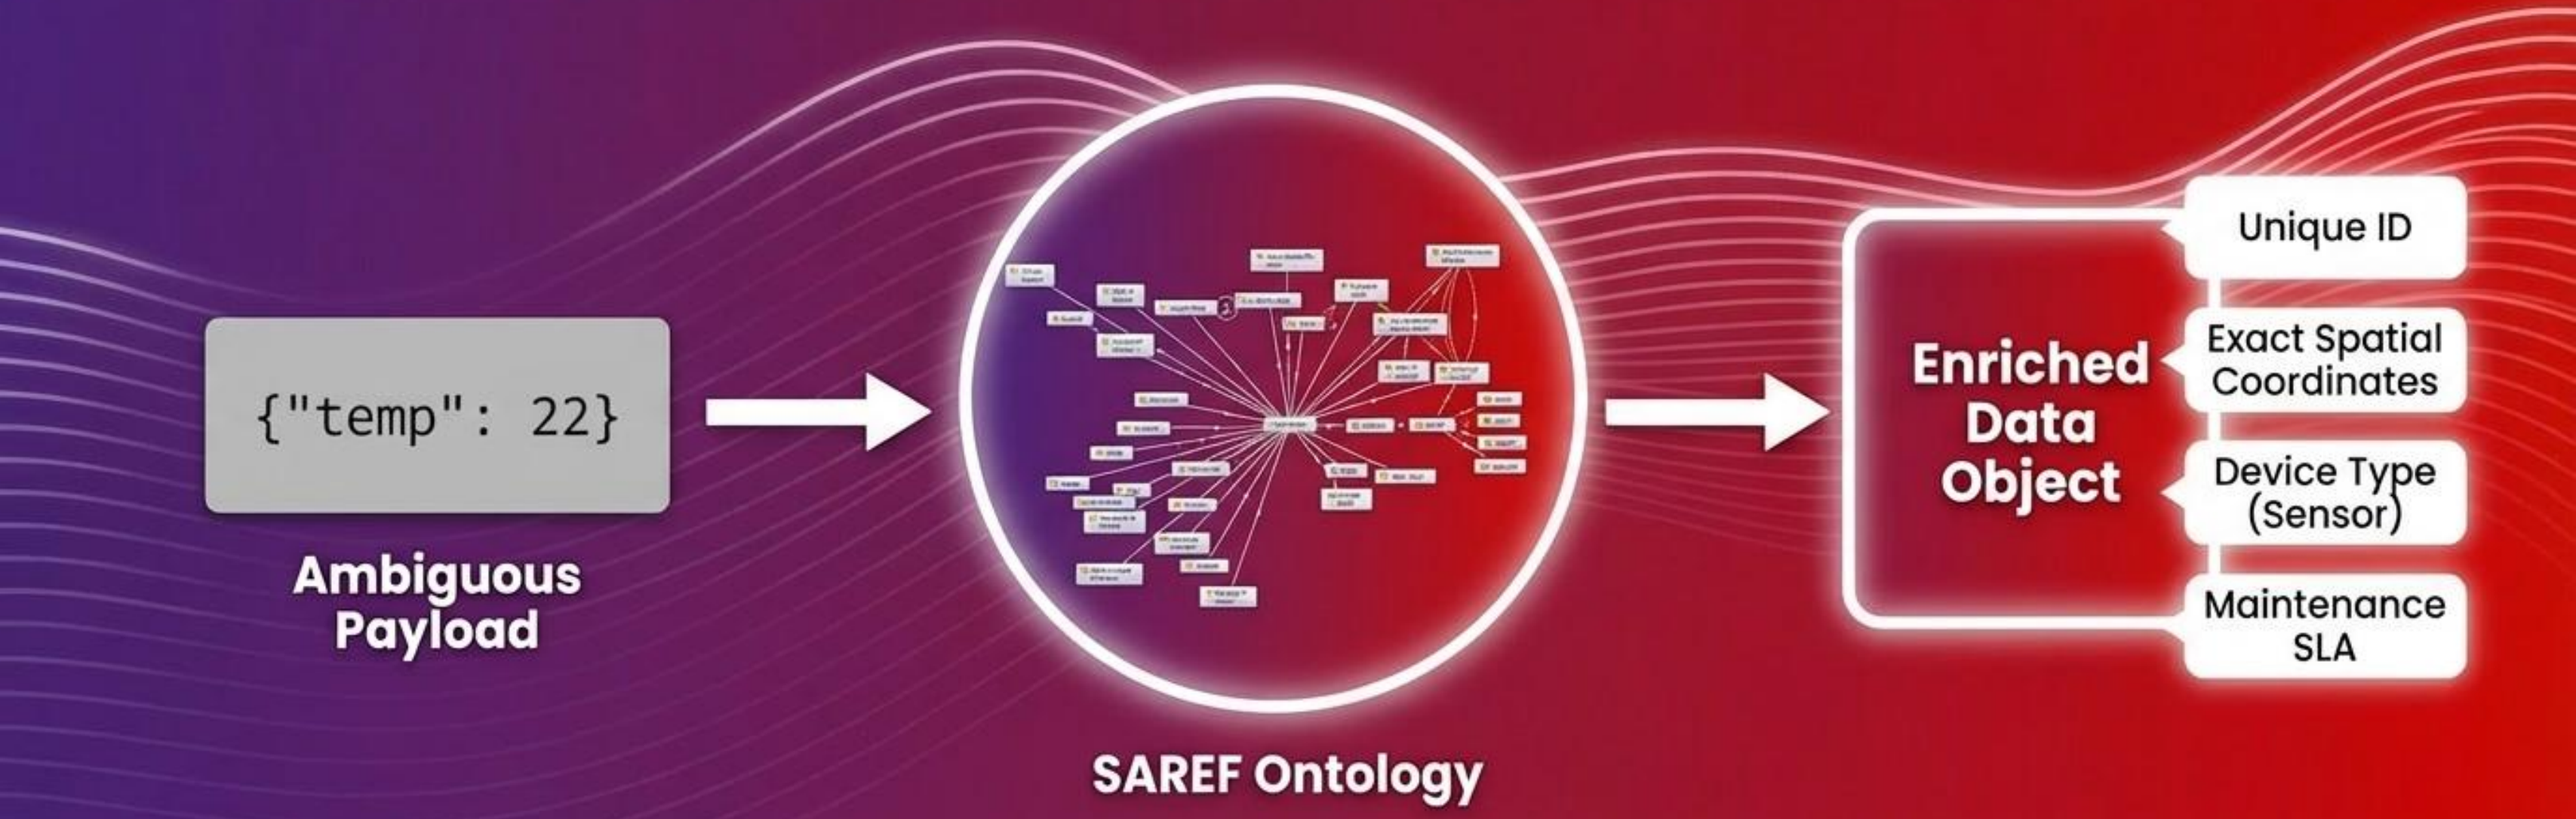

- Contextual routing ensures telemetry is instantly bound to its precise geographical and operational reality.
- Prevents critical misunderstandings by enforcing semantic consistency across any organization adopting the data.

The Transformation Funnel

Elevating raw street-level data into city-wide knowledge.

Digital Twin

The real-time, functional 3D replica.

City Digital Twin Data Model

The harmonized UrbanOS semantic layer.

Service Point

The geographic and operational context layer.

Data Object

The raw, standardized telemetry payload.

Device

The physical hardware on the street.

From Static Blueprint to Living Organism

The continuous operational loop of the UrbanOS.

1. Physical Asset Deployment

2. Real-Time Telemetry Generation

3. UrbanOS Data Curation & Contextualization

4. Digital Twin Simulation & Monitoring

5. Optimized Physical Maintenance

The Asset Lifecycle Loop

- Integrates real-time events and enables simulation scenarios based on synthetic, contextualized information.
- Data curation is a continuous, automated process, not a batch-job afterthought.

The Global Standards Architecture

The exact protocols enabling interoperability at scale

Category	Standard	Architectural Contribution
Identification	ISO/IEC 15459, RFC 8141 (URNs), Plus Codes	Unique, persistent, decentralized geospatial IDs.
Semantics	ETSI NGSI-LD, ISO/IEC 11179 (Smart Data Models)	The interoperability pillar; semantic governance.
IoT Protocols	OMA LwM2M, MQTTS/Sparkplug B	Deterministic mapping for constrained edge devices.

Deterministic Identification

The anatomy of a persistent digital identity

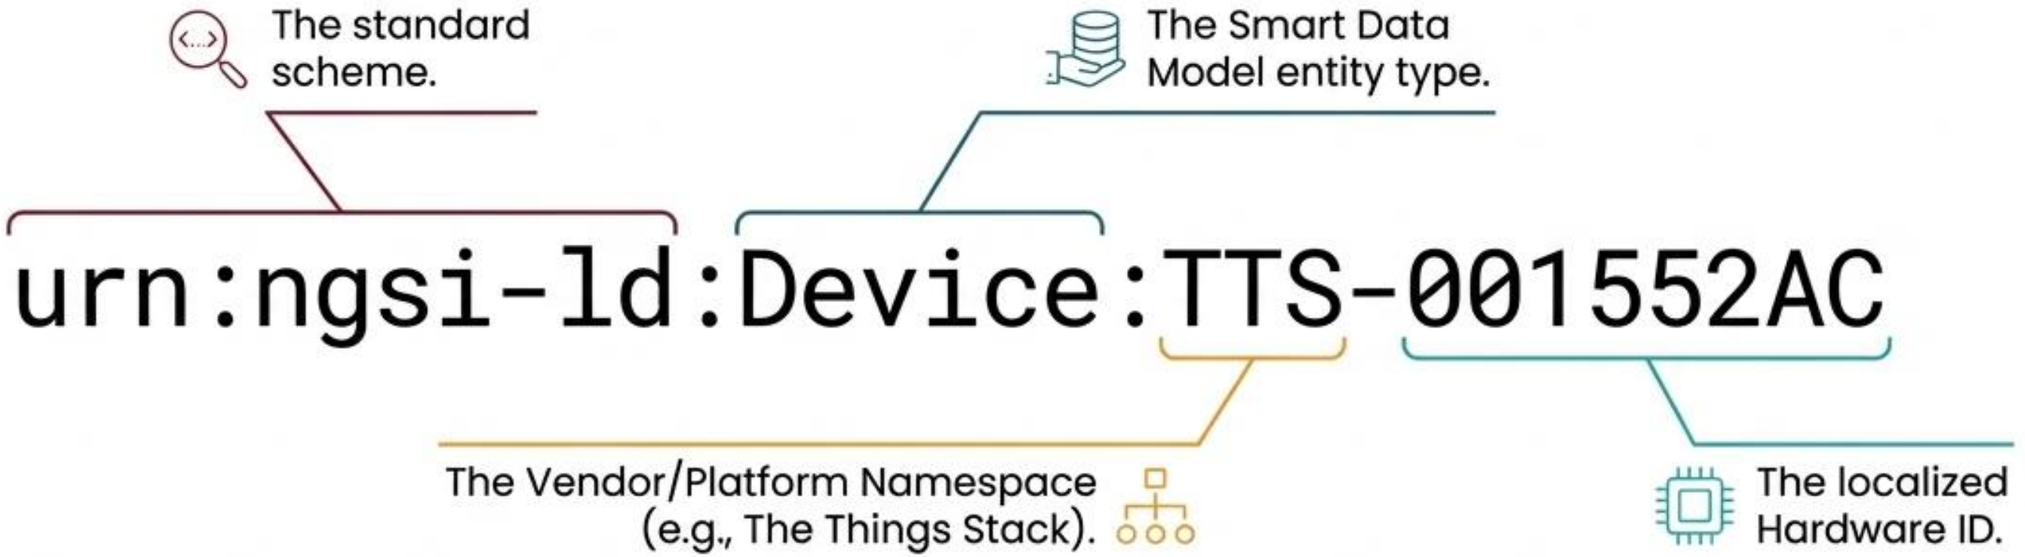

- Namespacing prevents collisions between heterogeneous systems (e.g., differentiating LoRaWAN nodes from proprietary platforms).

Pillar 1: Layered Spatial Architecture

Decoupling the physical world from the digital sensor

Top Layer (The Anchor):
ServicePoint

Middle Layer (The Purpose):
ServiceContract

Bottom Layer (The Canvas):
SUS (Smart Urban Space) | Example: urn:ngsi-ld:SUS:Valdemingomez

- Hardware is volatile; locations are permanent. This architecture ensures historical integrity even when devices fail.

The ServicePoint: A Virtual Anchor

Algorithmic geographic identifiers via Plus Codes

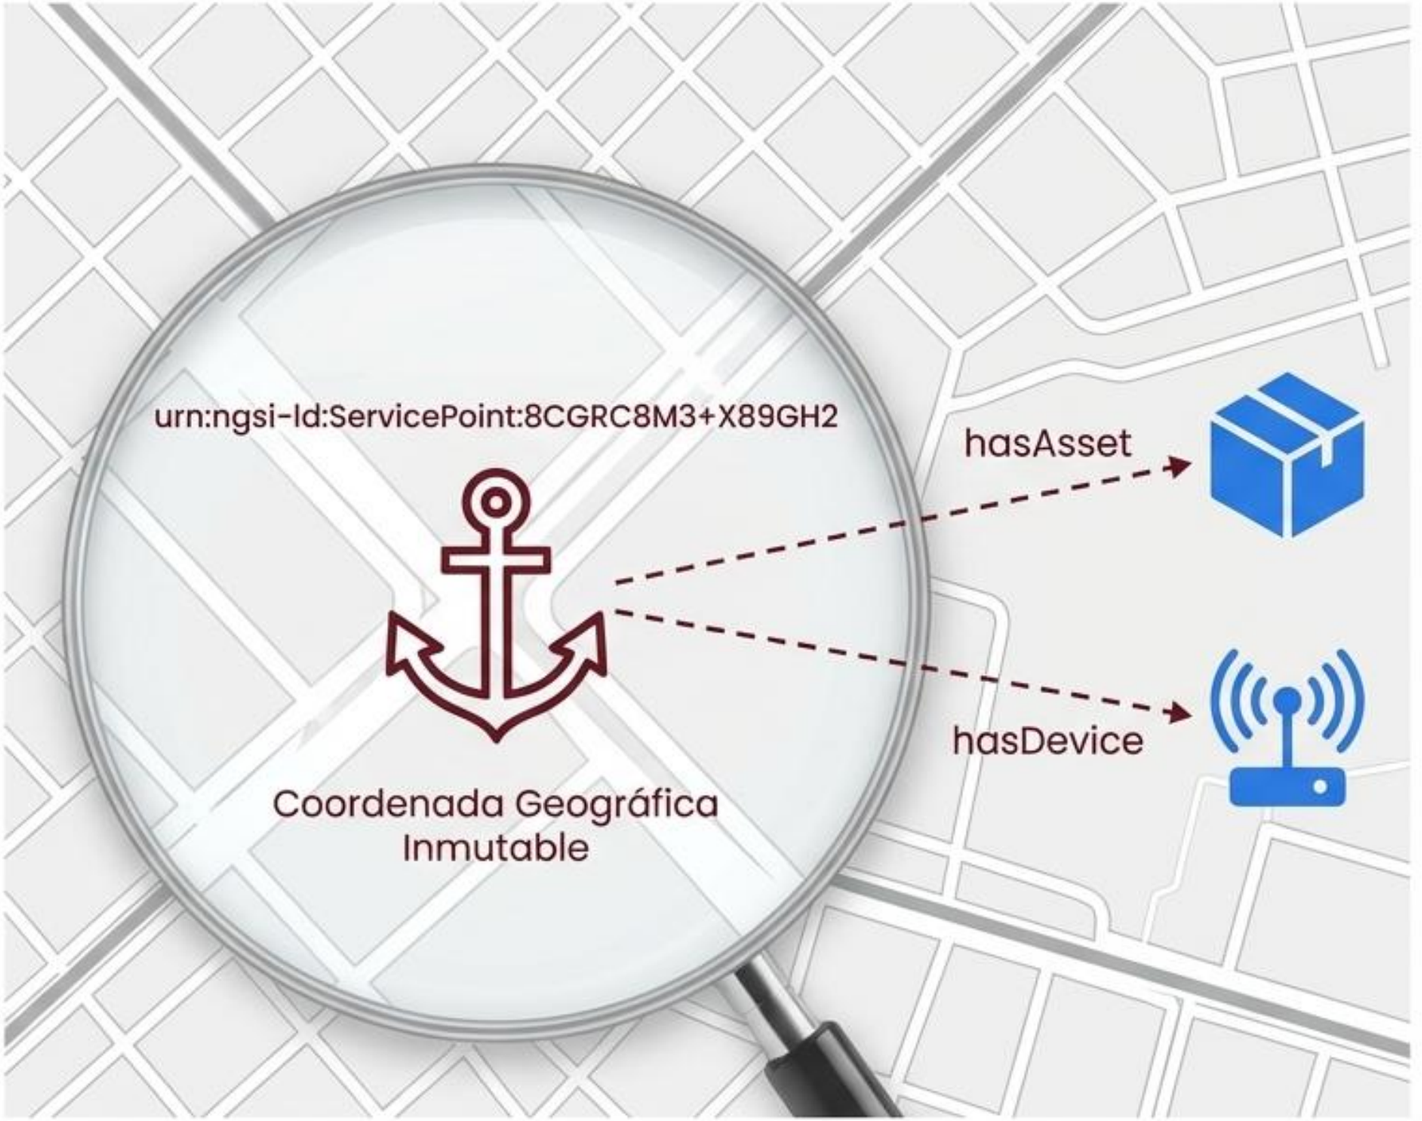

- The URN is not a sequential number, but a 15-character algorithmic hash of its exact coordinates, guaranteeing centimetric precision without a central authority.

Reconstructing History: Data Stitching

Preserving the service history, not just the device history

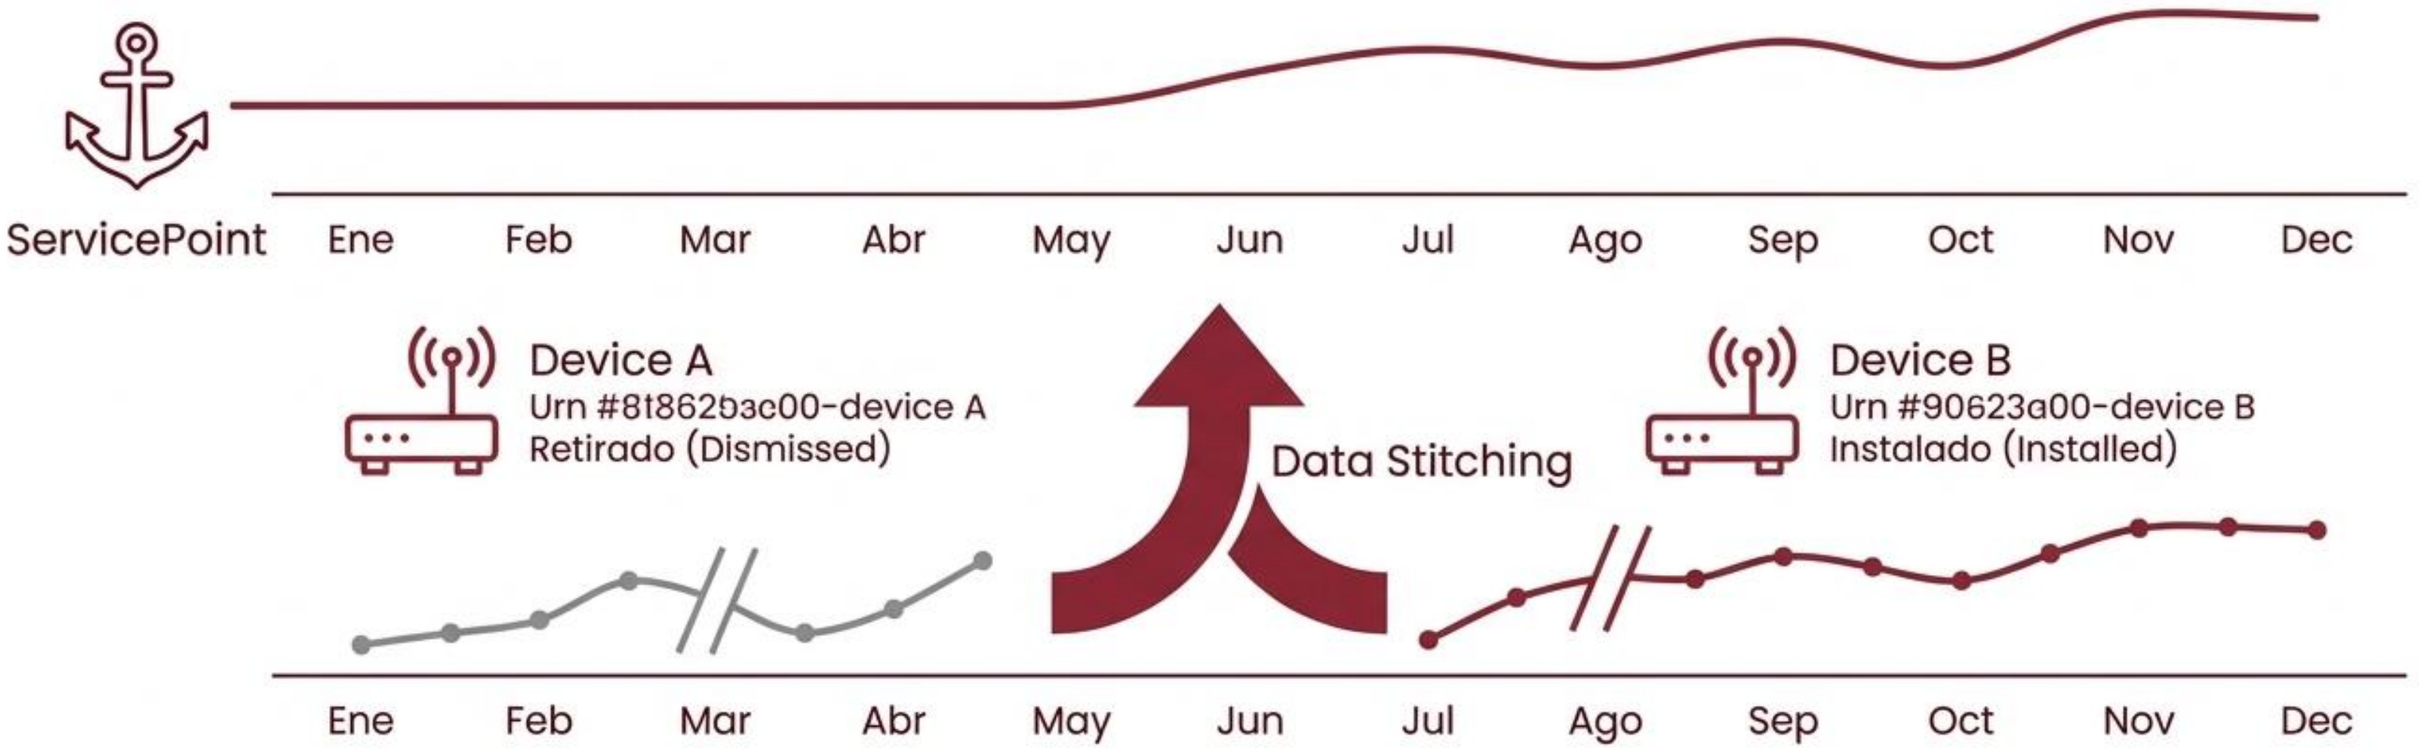

- Querying the ServicePoint seamlessly stitches telemetry from every sensor that has ever occupied that specific location.

Pillar 2: Asset Lifecycle Management

Modelling the physical reality through finite states

lifeCycleState
(The logistical phase)

lifeCycleSubState
(The contextual reason)

- Every physical asset is governed by a strict dual-state machine to ensure data quality and operational traceability.

The Dual-State Flowchart

Granular tracking from deployment to decommissioning

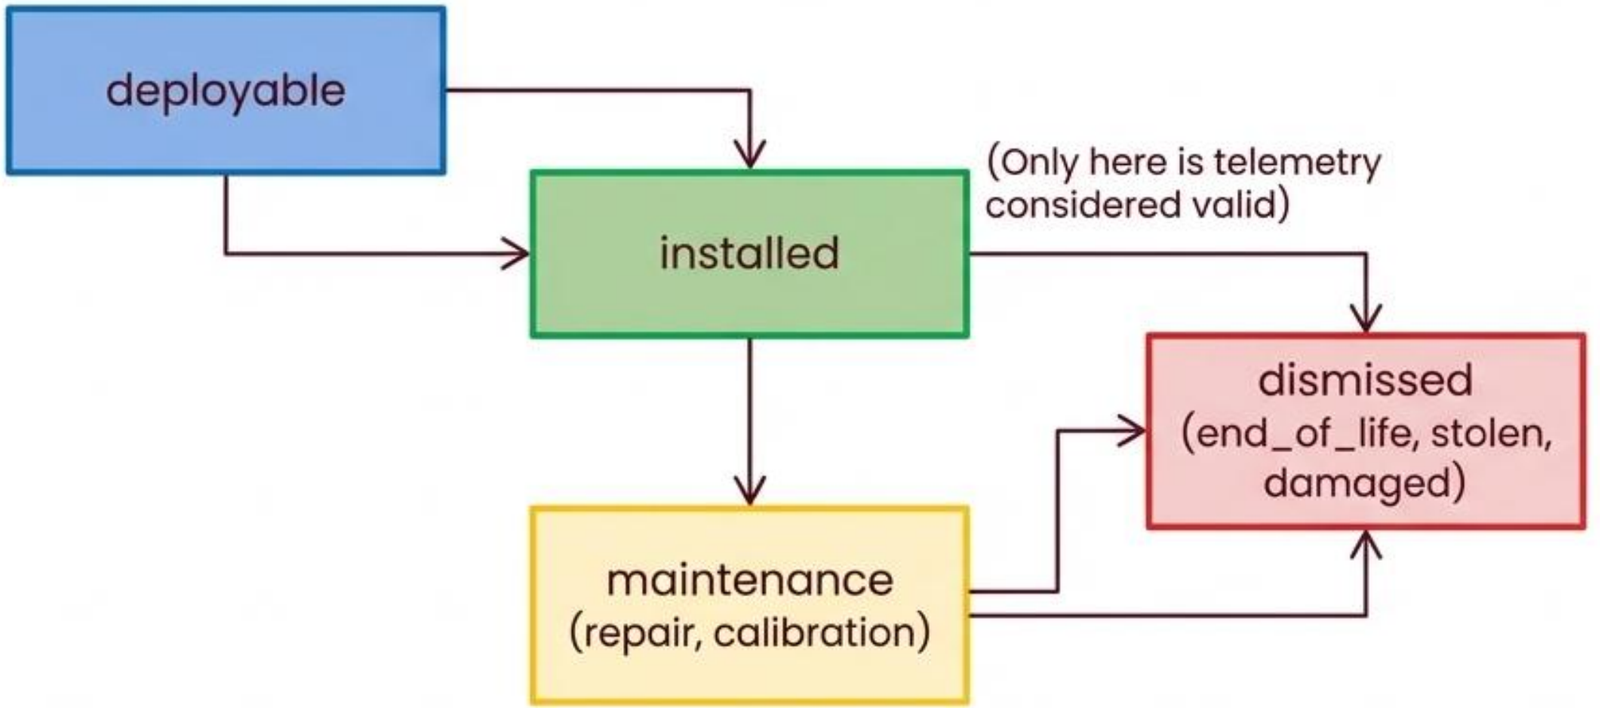

- Analytics algorithms filter out noise by only trusting telemetry from assets explicitly marked as ‘installed’.

Total Traceability: Snapshot & Log

The immutable audit trail for SLAs

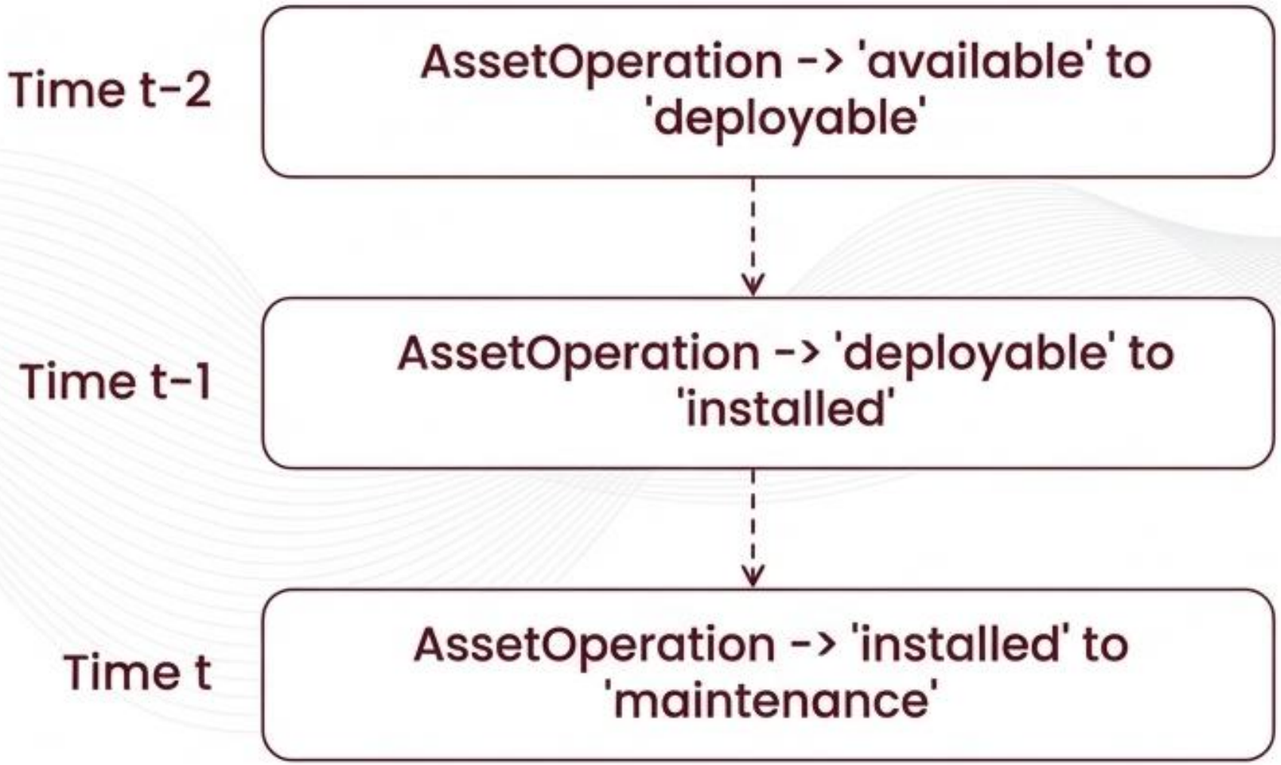

- Provides a fixed snapshot of the current state while preserving an auditable ‘movie’ of all interventions to calculate **precise MTTR** (Mean Time To Repair).

Pillar 3: The Semantic Stack

Guaranteeing that heterogeneous systems speak the same language

- Transforming a raw numeric value into interoperable, logically defined knowledge.

The Semantic Rosetta Stone

Deconstructing an NGSI-LD Payload

- This integration quietly converts a raw float ('0.75') into formal logic: 'The container is at 75% capacity according to the ETSI ontology.'

The Power of Graph Traversal

Answering complex operational questions in milliseconds

The Query: 'A sensor failed. Is it under warranty? Which company do I notify?'

The Result: Automatic identification of the responsible vendor and contract validity without manual database cross-referencing.

Behind the change. Beyond the challenge.

The Master Asset Catalog

A semantic firewall against vendor lock-in

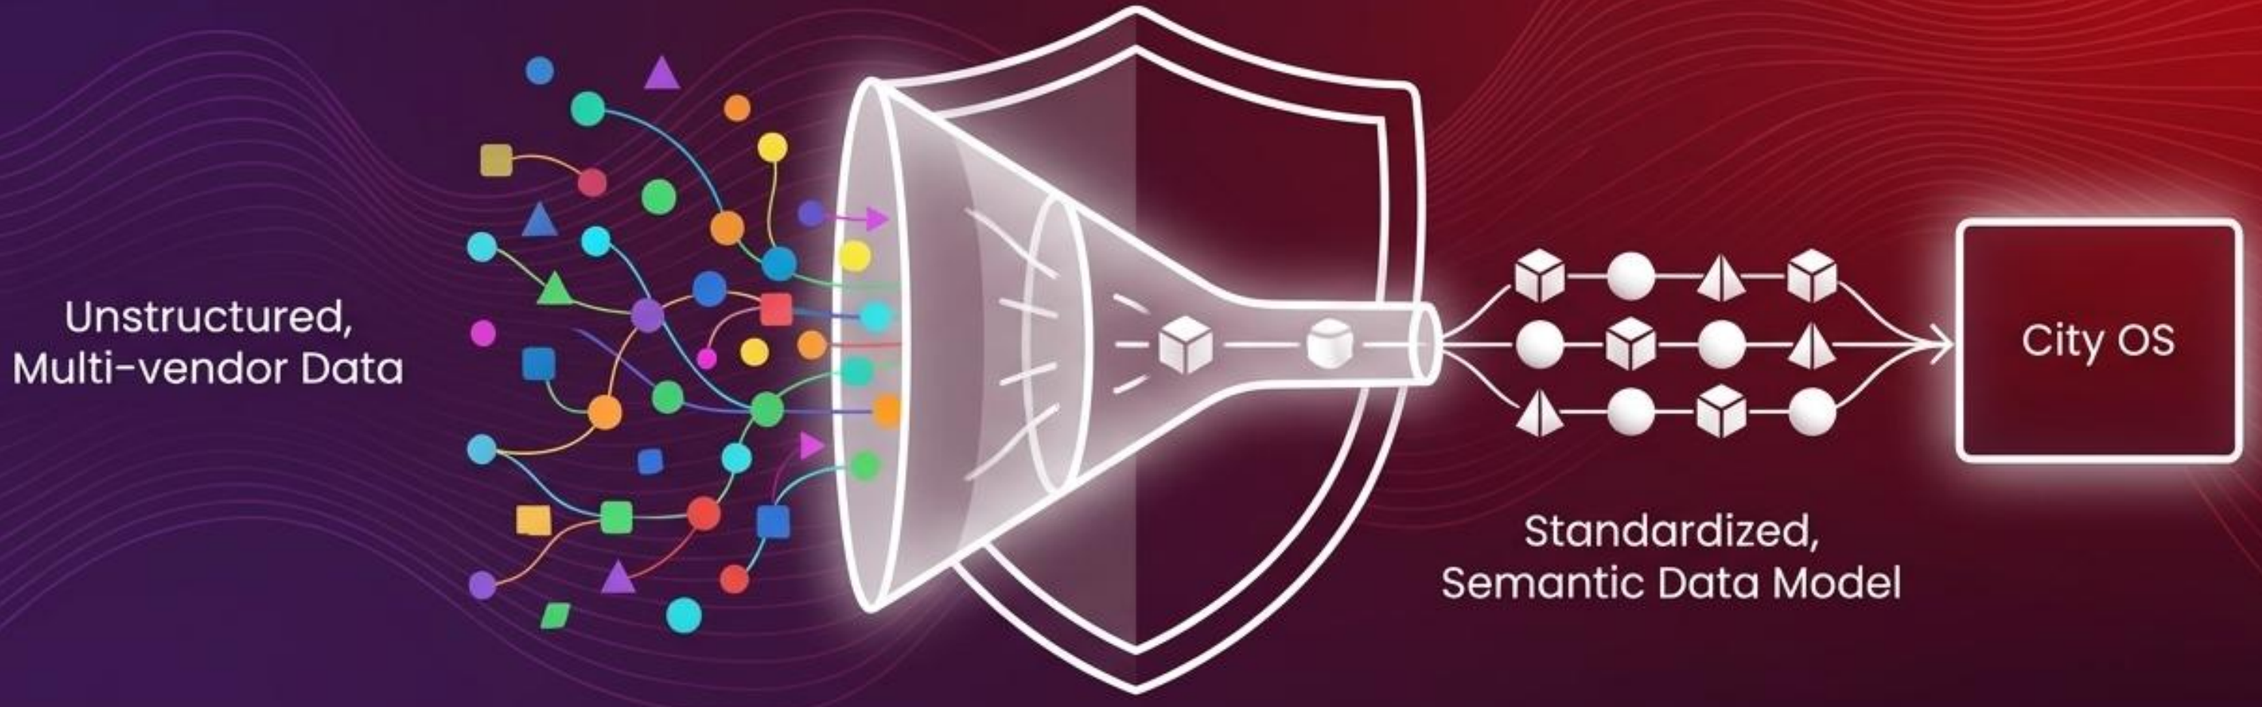

- Homologation via the refModel attribute ensures all new acquisitions are natively compatible with the city's architecture.

Enabling the Future

The systemic impact of semantic interoperability

Data Integrity

A single, auditable source of truth.

Real Interoperability

Breaking vertical silos; plug-and-play provider integration.

Smart Asset Mgmt.

Total lifecycle visibility for optimizing SLAs.

Digital Twin Cornerstone

The relational foundation for the city's virtual mirror.

Behind the change. Beyond the challenge.

Madrid, Digital Capital

Cementing the future on reliable data.

- By conquering semantic interoperability, Madrid has built the foundation for the 'Citiest 5.0'.
- Because Digital is Capital.

libelium eui | MADRID

Behind the change. Beyond the challenge.