iOM

WHITEPAPER

FOUNDATIONS · AI GOVERNANCE

Representation Is Not Authority

Infrastructure Operating Models vs. Digital Twins — why representation and authority are different functions, and legitimacy is the boundary.

ABSTRACT

As enterprises adopt automation and AI-driven operations, two concepts are routinely conflated: digital twins and Infrastructure Operating Models. They are not variations on one idea — they perform different functions. A digital twin **represents** infrastructure: it mirrors observed state to improve visibility, simulation, and explanation. An IOM **governs** infrastructure: it establishes authoritative intent, ownership, and legitimacy before anything is created or changed. A twin can tell you a resource exists, that it changed, that its traffic increased. Only an IOM can tell you that the resource should not exist. That judgment — legitimacy — is the real boundary of the category, and it is the one thing representation can never produce.

theIOM.org · Infrastructure Operating Model

Proposed vendor-neutral standard · v1.0 · CC BY 4.0

1. The purpose of a digital twin

A digital twin is designed to **describe reality**. In infrastructure contexts it mirrors assets, configurations, and relationships; reflects observed runtime state and telemetry; supports visualization, simulation, and “what-if” analysis; and helps operators understand behavior and dependencies.

The core value of a twin is **fidelity of representation** — the closer the model mirrors reality, the more useful it is for explanation and analysis. Twins are therefore anchored in runtime observation: when infrastructure changes, the twin updates to reflect what exists. That design is powerful, but it also defines the ceiling of what a twin can do.

2. The purpose of an Infrastructure Operating Model

An Infrastructure Operating Model is designed to **govern reality**, not merely reflect it. An IOM introduces an authority layer that defines what infrastructure is allowed to exist; how it may behave, interact, and evolve; who owns outcomes and escalation paths; which changes are legitimate before execution; and how observed behavior is reconciled against governing intent.

Rather than mirroring reality, an IOM establishes **normative constraints** that reality must conform to. Its value lies not in visibility, but in decision authority.

> **THE CATEGORY BOUNDARY**
>
> A digital twin can tell you a VM **exists**, that it **changed**, that its **traffic increased**. An IOM can tell you something a twin never can: that the VM **should not exist**. The first three are observations. The last is a judgment of legitimacy — and legitimacy, not representation, is where the category begins.

3. What “intent” means in an IOM

Intent is the explicit, authoritative declaration of how infrastructure is allowed to behave, interact, and evolve — independent of how it is implemented. Intent answers *“what must be true,”* not *“how it is done.”* That separation is everything.

What intent is not

Intent is not a configuration file, a Terraform plan, a firewall rule, a ticket description, a policy checkbox, or a diagram in a wiki. Those are implementations. Intent is the governing meaning behind them — and it exists as a separate concept so that meaning survives change. Tools, vendors, automation, and AI agents all change; if meaning is embedded inside scripts, configs, and tickets, it is fragile and easily lost.

4. The three-layer model

Infrastructure reality always exists across three layers. An IOM governs by continuously reconciling all three; a digital twin primarily models the third.

Layer	What it answers	Character
1 · Intent	Why / what must be true — e.g. “customer data must never leave region X”	Stable, human-meaningful, governing
2 · Implementation	How it is done — subnets, security groups, IAM, routing, Terraform	Flexible, replaceable, tool-specific
3 · Runtime reality	**What is actually happening — live traffic, active permissions, drift**	**Dynamic, often surprising**

Digital twins model Layer 3. IOMs govern Layer 1 and constrain Layers 2 and 3.

5. Types of intent

Intent is concrete, not abstract. An IOM distinguishes several kinds:

- **Architectural intent** — allowed structures: “production and non-production must never share identity boundaries.”
- **Policy intent** — why a rule exists: the rule is “port 22 blocked”; the intent is “no unmanaged administrative access paths.”
- **Ownership & accountability intent** — “this team owns blast radius for this dependency chain.”
- **Operational intent** — how change is allowed: “this class of changes may be autonomous; this always escalates to a human.” This is how governed autonomy works.

6. Where intent lives today

In most enterprises intent exists — but implicitly. It is scattered across tickets and approvals, architecture documents, policy statements, IaC repositories, and tribal knowledge. Within ITSM models intent is implicit, reconstructed after the fact, validated by humans, and lost once execution completes. ITSM coordinates workflows; it does not govern architectural legitimacy.

7. Representation vs. authority

The core distinction between a digital twin and an IOM is the difference between representation and authority.

Dimension	Digital twin	IOM
Primary function	Describe infrastructure	Govern infrastructure
Source of truth	Observed state	Declared intent
Relationship to execution	Post-hoc	Pre-execution
Role of telemetry	Central	Secondary
Treatment of drift	Detected	Invalidated
Governance	External	Embedded
AI safety	Informational	Deterministic

A digital twin answers *“what exists and what might happen?”* An IOM answers *“is this allowed, and should it proceed?”*

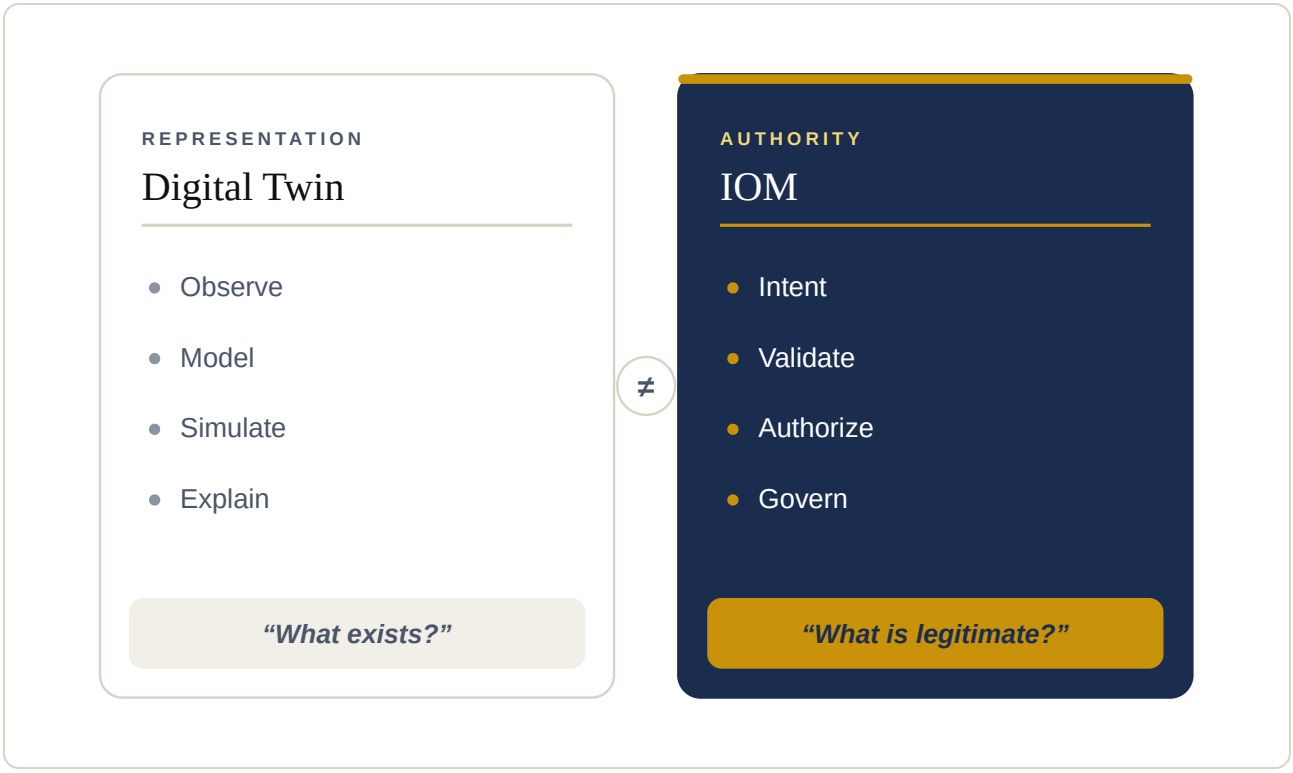

Representation and authority are different functions. A digital twin observes, models, simulates, and explains — answering “what exists?” An IOM declares intent, validates, authorizes, and governs — answering “what is legitimate?”

8. Why observation cannot create authority

Because twins are derived from observed state, meaning can only be inferred or annotated retrospectively. Intent is interpreted after execution, illegitimate states remain valid until corrected, and governance lives outside the model. AIOps shares the limitation: it reasons over telemetry to improve detection, correlation, and remediation, but it cannot determine whether an action is legitimate.

Observation cannot define intent

No volume of telemetry reconstructs why a system was built the way it was. Observation captures what is happening, never what was meant. Intent is a prior fact — a decision — and decisions are not recoverable from their effects.

Observation cannot establish legitimacy

The deeper limit is not about intent but about legitimacy. A twin can confirm that a connection exists and is active; it cannot say whether that connection is *allowed*. Legitimacy is not a measurable property of the running system — it is a relation between what is running and what was authorized. Observation has access to only one side of that relation. A twin can describe the world exhaustively and still be unable to declare a single thing illegitimate, because the standard of legitimacy was never in the telemetry to begin with.

> Authority is not a higher resolution of observation. It is a different kind of fact.

An IOM inverts the relationship by declaring intent and authority first, and requiring execution to conform before change occurs. AI may recommend; the IOM decides. Representation optimizes understanding. Authority establishes legitimacy.

9. From ticket-based to intent-based infrastructure

Ticket-based infrastructure exists because intent is implicit — humans must interpret meaning, assess risk, and approve legitimacy. Intent-based infrastructure replaces human arbitration with deterministic validation: changes are validated against intent before execution, tickets become exception handling, and control moves from humans to the operating model.

WHY SPEED-OF-LIGHT MATTERS

Execution operates at machine speed; control cannot operate at human speed. When execution approaches physical limits but control remains ticket-based, risk accumulates faster than it can be resolved. Ticket-based systems are bounded by human cognition; intent-based systems are bounded only by computation.

10. Why this distinction matters for AI

AI systems can reason over representations. They cannot derive authority from them. A model can ingest every signal a digital twin produces and still have no basis for deciding whether an action is legitimate — because legitimacy is not present in the representation; it must be supplied. This is why autonomous infrastructure stalls: given representation, AI can predict, explain, and recommend; only given authority can it act.

AI + digital twin reasons over observed state, simulates outcomes, and still requires human approval to determine legitimacy. **AI + IOM** reasons over declared intent, validates actions before execution, and enables safe autonomy. Digital twins make AI smarter; the IOM makes AI governable.

> Authority must exist before AI can safely act. Representation, however complete, is not authority.

Conclusion

Digital twins help enterprises see; Infrastructure Operating Models help enterprises decide. But the deeper point of this paper is not that one tool is better than another. It is that representation and authority are **different functions**. Representation answers what exists. Authority answers what is legitimate. No increase in the fidelity of the first ever produces the second.

As automation and AI accelerate, representation alone becomes not just insufficient but unsafe. Legitimacy must be declared before execution, because it cannot be recovered after it.

> Desired state is a signal. Intent is the authority.

Appendix — How intent exists across operating models

Dimension	CMDB	ITSM	AIOps	IOM
Primary purpose	Asset inventory	Workflow coordination	Ops optimization	Infrastructure governance
How intent is represented	Not represented	Implied in tickets	Inferred statistically	Explicit, first-class
Location of intent	Human memory	Free text, docs	Learned models	Canonical model
Relationship to execution	Post-hoc	Post-hoc	Reactive	Pre-execution
Authority over change	None	Human arbitration	Advisory	Deterministic
Drift handling	Recorded	Investigated	Detected	Invalidated
AI suitability	**Informational**	**Human-gated**	**Probabilistic**	**Deterministic**