

SIX LEVELS OF INFRASTRUCTURE AUTHORITY

The Infrastructure Authority Maturity Model

A vendor-neutral way to locate any organization on the path from tool-driven guesswork to authoritative, AI-ready infrastructure — and to see what each climb requires.

THE SIX LEVELS

- | | |
|--|---|
|  Level 0 — Tool-Driven | Knowledge lives in tools and people. |
|  Level 1 — Documented | Someone wrote it down. |
|  Level 2 — Observable | You can see what is happening — after it happens. |
|  Level 3 — Modeled | A real, reconciled model of what exists. |
|  Level 4 — Governed | Change is validated before it executes. |
|  Level 5 — Authoritative | One authority every actor must clear. |

HOW TO READ THIS MODEL

The model measures a single thing, ascending: how authoritatively an organization can say **what exists, what is intended, and what is permissible** — and enforce it before change executes. It is deliberately vendor-neutral. You can place any organization on it, including your own, without reference to any product.

Most enterprises sit at **Level 2 (Observable)**: strong monitoring, no authority. That is not a failure — it is the top of the *Observe* ladder. The *Govern* ladder above it is a different climb, and it is the one that determines whether automation and AI can act safely.

One axis, not two. AI-readiness is not a separate track bolted onto the top. It is what authority unlocks: at the highest level, autonomous AI can act safely because it is structurally unable to act outside authorized state.

LEVEL 0 · TOOL-DRIVEN

Reactive — you respond to infrastructure

Knowledge lives in tools and people.

- Truth is scattered across a CMDB, spreadsheets, monitoring, and the heads of a few key engineers.
- There is no single source of truth — the answer depends on who you ask.
- Change is approved by ticket and tribal knowledge, then verified by hoping nothing broke.

Ceiling. Every change is a guess. The cost of a mistake scales with how fast you can execute — and nothing can see the mistake coming.

LEVEL 1 · DOCUMENTED

Someone wrote it down.

- Runbooks, architecture diagrams, and wikis exist.
- Documentation is written for humans and drifts out of date the moment infrastructure changes.
- Automation and AI cannot reliably operate against it.

Ceiling. Documentation describes the past. It is already stale, and no machine can trust it to make a decision.

LEVEL 2 · OBSERVABLE

You can see what is happening — after it happens.

- Monitoring, logging, and tracing give rich visibility (Datadog, Splunk, and the like).
- Teams detect and respond to problems faster than ever — a real achievement.
- Signals describe current behavior, not intended state or what is permitted.

Ceiling. Observability is reactive by design. It tells you a change broke something — after the change ran. It cannot validate a change before it executes.

LEVEL 3 · MODELED

Aware — you understand infrastructure

A real, reconciled model of what exists.

- A continuously reconciled model captures assets, dependencies, and ownership.
- Teams can answer “what do we have, and what depends on it?” with confidence.
- The model describes reality — but does not yet govern what happens to it.

Ceiling. Most organizations that reach here believe it is the finish line. It is not. Knowing the environment is not the same as governing change to it.

LEVEL 4 · GOVERNED

In authority — you govern infrastructure

Change is validated before it executes.

- Intent and policy are encoded in the model, not just documented.
- Every proposed change is checked against ownership, dependencies, and policy before it runs.
- Non-compliant changes are denied with a reason; approved changes carry an audit trail by default.

Ceiling. Governance now covers human and automated change. The final step is extending the same authority to autonomous actors that move faster than any reviewer.

LEVEL 5 · AUTHORITATIVE

One authority every actor must clear.

- Humans, automation, and AI agents all validate against the same authoritative model before touching infrastructure.
- An actor — human or AI — can only propose. The model validates and the engine executes; nothing acts on infrastructure unchecked.
- Compliance is continuous, and every action is admissible by construction.

The payoff. The top of the ladder: the state in which autonomous AI can act safely — because it is structurally unable to act outside authorized state.

HOW TO USE THIS MODEL

1. Locate yourself. Read the levels and find the highest one you can honestly claim across your whole estate — not your best team on its best day.

2. Name the ceiling you are hitting. Each level fails in a specific way. The ceiling text tells you what your current level cannot do, no matter how well you execute it.

3. Scope the next climb. Moving up is a defined step, not a leap:

- **0 → 1:** Consolidate scattered knowledge into one place that is written down and shared.
- **1 → 2:** Move from static documents to live signal about what is actually running.
- **2 → 3:** Add a reconciled model of what exists and what depends on what — not just what is happening.
- **3 → 4:** Encode intent and policy so the model can validate change before it executes.
- **4 → 5:** Make the model the single authority every actor — including AI — must clear.

The honest finish line. Levels 0–2 are reactive — you respond to infrastructure. Level 3 is aware — you understand it. Levels 4–5 are authoritative — you govern it. Only at the top can autonomous AI act safely, because the authority to act is enforced by architecture, not by oversight you hope holds.

WHERE TO START

The fastest way to place your organization is to model a slice of your real environment. An Infrastructure Authority Assessment returns your level and the scoped gap to the next one — against your actual estate, not a questionnaire. Start at authoriom.com/assessment or contact sales@authoriom.ai.