What Is Decay Maturity? The Seventh Dimension of Organizational Metabolism

By Rahul Jindal

Six months after the AI feature shipped, retention has slid four points and nobody on the AI council has connected the two. The system is up. The pipelines are green. The model still serves traffic. And it is quietly worse than it was on launch day. This is the gap Decay Maturity measures.

Decay Maturity is the seventh dimension of the Organizational Metabolism Index, available in OMI Enhanced. The first six dimensions (Leadership, Process, Talent, Data, Technology, Culture) measure how fast your enterprise can absorb AI. Decay Maturity measures something the first six do not: whether what you absorbed is still working.

> ***"Capability tells you what your AI can do on its best day. Decay Maturity tells you what it actually does on the average day, six months in."***

Why a seventh dimension was needed

The original six dimensions of OMI describe absorption: how fast capability moves from possibility to operational reality. That was the right frame for the era when the binding constraint was getting AI deployed at all. By 2026, most enterprises have crossed that threshold. The question shifts. What used to be hard was getting the thing in. What is hard now is keeping the thing alive.

MIT's Project NANDA, RAND, and Gartner each reported in 2025-2026 that between eighty and ninety-five percent of enterprise AI initiatives produce zero measurable ROI or have been abandoned. Most of the surviving systems are quieter than that. They still run, just less well than the day they shipped. The gap between launch-day capability and six-

months-in capability is what we call the Decay Tax. Decay Maturity is the dimension that measures whether your enterprise is set up to resist it.

The eight currencies of decay

Decay is not one phenomenon. It is at least eight, and most enterprises track none of them.

1. **Model drift.** The input distribution shifts; accuracy drops. The classical case.
2. **Eval debt.** The gap between what you tested and what users actually ask now.
3. **Silent regression.** The vendor swaps the model under you; behavior changes; nobody pages.
4. **Prompt rot.** The prompt that built the demo does not survive two model versions.
5. **Agent error compounding.** A 95% reliable step run five times in series gives you 77% reliability. Run it ten times: 60%.
6. **Hallucination accumulation.** Each false answer is small; across millions of interactions the trust hit compounds.
7. **Refusal and capability drift.** Yesterday's working prompt is today's refused request; yesterday's well-bounded answer is today's overconfident essay.
8. **Shadow AI.** Employees use the consumer stack outside your governance, with data you cannot see.

Decay Maturity asks how well your enterprise instruments, owns, and acts on these eight currencies, without requiring that you have solved them all. It is a maturity dimension, not a perfection one.

What the dimension actually measures

The five questions that compose the Decay Maturity dimension are deliberately blunt. They probe what an honest engineering and operations leadership team should be able to answer without preparation.

- Do you measure how AI accuracy or quality changes over time after a system is deployed?

- When a vendor swaps the underlying AI model, how do you catch behavior changes?
- Do your AI systems have a measurable half-life, the period after which 50% of accuracy or utility has eroded?
- How quickly can you detect that an AI system in production is degrading?
- Who owns AI reliability and decay across your enterprise?

Five questions, five-point scale each, average mapped onto a 0-100 dimension score, contributing to the overall OMI Enhanced score. The score then maps to one of four stages.

The four-stage maturity model

Decay Maturity is a stairway, not a switch. Most enterprises are on the first or second step. The point of the model is to name where you are honestly so the next move is obvious.

Stage 1

Blind

Most enterprises
What it looks like: AI is in production. You measure capability and cost. You do not measure decay. You will discover the Decay Tax through a customer escalation or a viral failure.

Next move: pick the single most business-critical AI system in production. Measure baseline accuracy this week. Re-measure in thirty days. The delta is your first data point, and the conversation starter your board will not be able to ignore.

Stage 2

Aware

Forward enterprises
What it looks like: you have a postmortem culture. You know drift exists. You have not yet instrumented it. Your AI council talks about it; nobody owns it. Conversation rich, accountability thin.

Next move: appoint a host. The Eighth Conversation (reliability) needs a named owner before it needs a tool. Read the case for a Chief Reliability Officer for AI for the structural argument.

Stage 3

Instrumented

Industry leaders
What it looks like: you measure half-life. Eval sets get refreshed on a cadence. You catch silent regression on vendor model swaps within hours, not quarters. You have a number to show the board.

Next move: publish your Half-Life Index. Set error budgets per system. Establish the runbook for a meaningful AI incident. The leap from Instrumented to Self-Healing is the leap from monitoring to acting.

Stage 4

Self-Healing

Frontier
What it looks like: your systems re-evaluate continuously, route around degradation, alert before users notice. Decay is a managed cost, not a hidden tax. Half-life is a board metric.

Next move: share the playbook. Become the industry benchmark. Your competitive advantage is now the speed at which you detect and route around erosion that competitors discover only after their customers do.

How to read your Decay Maturity score

Inside the OMI Enhanced report, your Decay Maturity dimension score appears alongside the other six on the same 0-100 scale, with the same Strong / Developing / Needs Attention labels. The translation back to the four stages is roughly:

- **0-25**: Blind. The instrumentation gap is the headline.
- **26-50**: Aware. You have started; the missing piece is the named owner.

- **51-75**: Instrumented. The work now is hardening, not building.
- **76-100**: Self-Healing. The work is leadership, externalization, and the next dimension nobody has named yet.

If Decay Maturity is your weakest dimension, that result is actually clarifying. Most enterprises have a hidden reliability deficit dragging on the rest of their AI program, and seeing it on the same chart as Leadership and Process is what gets it on the steering committee agenda.

Where Decay Maturity sits in the OMI family

The three OMI tiers compose like Russian dolls. Each tier uses the same scoring engine; each adds depth without changing the engine.

- OMI Lite : Leadership, Process, Talent. Three minutes. The fastest read on absorption maturity.
- OMI : adds Data, Technology, Culture. Five minutes. The canonical diagnostic, benchmarked against twelve industries.
- OMI Enhanced : adds Decay Maturity. Seven minutes. The full diagnostic plus the back-half-of-the-lifecycle dimension.

Enhanced is also where the next dimensions land first as the framework grows. Reliability is the dimension we needed today. Agent governance, vendor portability, and others are already in conversation.

What to do with this

If your enterprise has any AI in production today, the value of the Decay Maturity dimension is the question it forces. Take it on yourself first. Take it on each AI system second. Whichever stage you land in, the next move is named on this page. The cheapest first action is not a tool. It is the question: *for each AI system in production, what is its half-life, who owns measuring it, and when did we last refresh the eval?* If nobody can answer for any system on the list, you are at Stage 1, and you have just found the most expensive question in your enterprise.

The Decay Maturity dimension is the operational landing for the Decay Tax framework: the umbrella concept that every AI system pays a tax in reliability, accuracy and trust as it ages. The Eighth Conversation essay argues for the role that should host it. To measure your own, take OMI Enhanced.
