AI GOVERNANCE ASSESSMENT

Which AI tier does your governance support?

This self-assessment turns the argument of the book into questions an executive can answer from their own position, without convening a programme or examining team behaviour. The questions follow the order of the book: investment governance and team accountability (Part II), AI governance (Part III), and the transition itself (Part IV). The answers reveal whether the governance conditions for a product operating model are in place, or whether the structure merely appears to be.

Read each question. Answer it as the governance machinery actually operates, not as the documentation says it does. Where the answer runs against the product model, that is where the design work sits. Two columns sit beside each question. The Ch column names the chapter that makes the full argument behind it. The Commitment column shows which of the four structural commitments in Chapter 15 the question tests: 1, one point of accountability; 2, persistent teams and outcome funding; 3, consumption-speed finance; 4, from order-taker to order-maker. A dash marks a question that tests a condition the commitments depend on rather than a commitment itself. The final section reads the answers back against the four commitments.

ORGANISATION _____ COMPLETED BY _____ DATE _____

The four commitments (Chapter 15)

1 **One point of accountability.** A single, accountable individual has ongoing authority over what the product team works on and what gets released.

2 **Persistent teams, outcome funding.** The team that commits to an outcome also bears the consequences of pursuing it, funded against outcomes rather than stage gates.

3 **Consumption-speed finance.** Expenditure is reported against the evidence the portfolio review needs, at the cadence the work is consumed.

4 **From order-taker to order-maker.** The technology function solves the business's problems rather than collecting its requirements.

Using it with a leadership team

1. Each member of the leadership team answers on their own copy before the session.
2. In the room, spend the time on the questions where the answers differ, and find out which part of the organisation each answer describes.
3. Score the team's agreed answers on the last two pages to find your tier, and decide which commitment to work on first.

LAYER 1

Investment Governance

	QUESTION	CH	COMMITMENT	YES	NO	PARTLY
01	Does technology investment stay with a product for as long as it pursues its outcome?	4, 6	2	☐	☐	☐
02	Does the performance reporting you receive measure the value customers get from each product?	7	3	☐	☐	☐
03	Can the Product Manager move a value stream's capacity away from work the evidence no longer supports, without a separate committee decision?	5, 7	1	☐	☐	☐
04	When a product's evidence turns negative, can its funding be reduced or redirected with no more approval than continuing it would need?	6, 7	2	☐	☐	☐
05	When your last technology program closed, did the team that built it stay together?	4	2	☐	☐	☐
06	Can every product in your portfolio be stopped when the evidence says so, whatever its investment model?	6	2	☐	☐	☐
07	Does the Sprint Review show what the Increment cost alongside what it produced?	15	3	☐	☐	☐
08	When Sprint evidence shows a product should change direction, can the team redirect its spending without waiting for the next budget cycle or a re-approval against a fixed scope?	15	3	☐	☐	☐
09	Is AI consumption reported for each product at Sprint cadence, to the person who holds the release decision?	6, 15	3	☐	☐	☐

LAYER 2

Team Accountability and Decision Rights

	QUESTION	CH	COMMITMENT	YES	NO	PARTLY
10	For each product in your portfolio, does a single, accountable owner answer for what it delivers, with the authority to act without a committee decision?	5, 15	1	☐	☐	☐
11	When a Product Manager stops a feature the business wanted, does the organisation accept it as their decision to make?	13	1	☐	☐	☐
12	Is each product team cross-functional, carrying discovery, development and operations itself under one persistent accountability?	5	2	☐	☐	☐
13	Can you trace a recent product outcome, such as a customer result or a value realised or missed, to a decision made by a single, accountable owner?	5	1	☐	☐	☐
14	Does each value stream have both escalation paths in place: value questions from the Product Owner to the Product Manager, and operating model questions from the Scrum Master to the Product COO?	5	1	☐	☐	☐
15	Have you set the limits within which product teams make decisions for themselves?	8	1, 2	☐	☐	☐
16	When a high-stakes product decision was last made below your level, did you let it stand as theirs to make?	8	4	☐	☐	☐
17	Do your teams receive customer problems to solve, instead of specifications to implement?	15	4	☐	☐	☐
18	Do customer needs and governance requirements such as security, compliance and audit reach the team through separate channels, so it can tell a constraint from a demand?	15	4	☐	☐	☐
19	Has an executive said, explicitly and publicly, that the technology function's job is to solve the business's problems rather than to collect its requirements?	15	4	☐	☐	☐

LAYER 3

AI Governance Readiness

	QUESTION	CH	COMMITMENT	YES	NO	PARTLY
20	For any AI system your organisation has deployed, can you name the governing instructions that specify what it is authorised to do without human review at each step?	9	–	☐	☐	☐
21	Are those governing instructions under version control, with a revision history, and audited on a regular cadence?	9, 11	–	☐	☐	☐
22	Does every skill file and agent instruction still follow the Definition of Done?	11	–	☐	☐	☐
23	Has your organisation decided, by design, where an AI system's output may lead to a consequential decision without a person reviewing the reasoning?	8, 9	1	☐	☐	☐
24	For each AI deployment, have you separated decisions that can be reversed within an iteration from those that cannot, and set human oversight to match?	8	–	☐	☐	☐
25	If an AI system had been producing incorrect outputs for three months, would a check your organisation runs today have caught it?	11	–	☐	☐	☐
26	Do Scrum Masters raise operating model problems with the Product COO, and does the Product COO act on them?	11	–	☐	☐	☐
27	Have your teams stayed together through the last budget cycle?	11	2	☐	☐	☐
28	Is AI team performance inspected and acted on at a fixed cadence through the product's own reviews, such as the Sprint Review?	10	–	☐	☐	☐

LAYER 4

Transition Progress

	QUESTION	CH	COMMITMENT	YES	NO	PARTLY
29	Twelve months into your transition, has the way investment is approved and reported changed, beyond the vocabulary?	6, 12	–	☐	☐	☐
30	Are the people who raised structural questions before the transition now in roles where those questions are welcome?	App. D	–	☐	☐	☐
31	When your organisation faces budget pressure or a political escalation, does it use the product model's mechanisms to respond?	12	–	☐	☐	☐
32	Is there a product where the senior leader now says, unprompted, "The Product Manager owns this decision"?	13	1	☐	☐	☐
33	For middle managers whose routine work is being automated, does a program that develops the judgement their roles will need start before that work is automated?	14	–	☐	☐	☐
34	Does every investment case that automates developmental work record the four fields Chapter 14 sets out: classification, replacement mechanism, accountability, and sequencing?	14	–	☐	☐	☐
35	Is your transition planned as ongoing change, with no completion condition that closes the program?	Intro, 15	–	☐	☐	☐
36	Has a lesson from last quarter's AI deployments changed how your organisation governs them?	15	–	☐	☐	☐

SCORING

Reading your results against the four commitments

COMMITMENT	QUESTIONS CARRYING IT	NO OR PARTLY	÷	SHARE
1 One point of accountability	03, 10, 11, 13, 14, 15, 23, 32		÷ 8	
2 Persistent teams, outcome funding	01, 04, 05, 06, 12, 15, 27		÷ 7	
3 Consumption-speed finance	02, 07, 08, 09		÷ 4	
4 From order-taker to order-maker	16, 17, 18, 19		÷ 4	

For each commitment, count the questions carrying its number that you answered No or Partly, and divide by the number of questions carrying that number. The commitments carry different numbers of questions, so compare these shares rather than raw counts. The commitment with the highest share is the one most likely to be missing in practice, regardless of what the organisation chart or transition plan says. Chapter 15 sets out the failure each missing commitment produces.

A clean result on three commitments does not offset a weak fourth. The properties the product operating model exists to produce—bounded autonomy, empirical inspection, and subsidiarity—appear only when all four commitments are in place. Partial commitment is how most transitions stall, and the failure that follows is the one the missing commitment existed to prevent.

Where the answers point to more than one weak commitment, start where Chapter 15 ends. Choose one product and fund it against an outcome rather than a set of deliverables; Chapter 13 explains why funding changes first. Make a single, accountable owner answer for that outcome and give them the release decision. Then attend the Sprint Review four weeks later and judge the working output against the outcome the money was committed to. Repeat this assessment after the first quarter of Sprint Reviews. The answers that have changed show which commitments are now operating. The answers that have not show where the design work still sits.

Questions adapted from Evolve: The Operating Model AI Demands (Hodgson, 2026), Appendix A. © 2026 Matthew Hodgson. Zen Ex Machina · zenexmachina.com/evolve/governance-diagnostic · Worksheet v1.2

YOUR TIER

Which AI tier your governance supports

STEP	COUNT	RESULT
1. AI governance. Questions 20, 21, 22, 23, 24, 25: how many did you answer Yes?		4 or more of 6: in place
2. The four commitments. From the scoring page: how many commitments have a share of 40% or less?		All 4: in place

TIER	WHEN	WHAT IT MEANS
1 **Personal tool**	AI governance not in place	People use AI one step at a time and stay in control of every step. Nobody has designed how it is governed, so the gains stay with the person using it: faster drafting, summaries and research, with no advantage the organisation can build on.
2 **Team accelerator**	AI governance in place; one or more commitments not in place	AI speeds up the workflows you already have, inside controls added to the existing operating model: policy, access rules and review gates. The governance is designed; the operating model around it is not. Returns rise, then stop at the ceiling of the process AI was added to.
3 **Enterprise capability**	AI governance and all four commitments in place	AI agents carry out multi-step work on their own, inside governance boundaries you have designed, and your processes are built around what they can do. Governance and the operating model were designed together, and each cycle revises the decision boundaries, so returns grow with every cycle.

Adapted from Boston Consulting Group (2025b).

The tier is the highest your governance supports, not a measure of how much AI you use. Chapter 1 sets out the three tiers. For what it takes to reach the next tier, enter the agreed answers at zenexmachina.com/evolve/governance-diagnostic for the full report.

Questions adapted from Evolve: The Operating Model AI Demands (Hodgson, 2026), Appendix A. © 2026 Matthew Hodgson. Zen Ex Machina · zenexmachina.com/evolve/governance-diagnostic · Worksheet v1.2