

ENTINOLOGY

Specialised AI-Driven Project Delivery

A RIBAIA FIELD GUIDE

Cape Town to Hawaii

Role-Based AI Adoption in Ten Steps

Roelof Vermeulen

Creator of Role-Based AI Adoption and EPM-AI

DEPARTURE LOUNGE

The jet at Cape Town International

Picture your company as an airline with one aircraft. The board bought a modern, fast and expensive jet. It is sitting on the apron at Cape Town International with its engines running.

Everyone has a boarding pass, but nobody has chosen a destination. The crew met at the gate. The catering manager holds the pilot's manual, while the pilot watches a video about reclining seats. Fuel burns as people photograph the cockpit.

Replace the jet with AI and you have a fair picture of many businesses in 2026.

The numbers support the comparison. In a 2026 YouGov survey of more than 500 US and UK business leaders, 59% reported an AI skills gap, although 68% said employees had access to AI learning resources. Access is not capability (DataCamp, 12 March 2026).

Spending is still climbing. A BCG survey of 2,360 executives across 16 markets found that companies expected to double AI spending in 2026. Leaders are buying more fuel for an aircraft that may still be parked (BCG AI Radar 2026, 15 January 2026).

I have spent nearly 30 years helping organisations deliver work. New tools often arrive before the operating method. Activity rises, while useful output remains hard to prove.

This book takes your company from Cape Town to Hawaii, roughly 18,500 km away with no scheduled nonstop flight. You do not land there by accident.

I created Role-Based AI Adoption, or RiBAIA (say ri-BY-ah), for that journey. Its ten steps turn a general ambition into named work, accountable roles, controlled outputs and measurable value.

You can read the book in about half an hour. Do not read it passively. Bring one real piece of work from your business and apply every step to it.

THE PROBLEM

One manual for the whole crew

AI often enters a business like this. Someone senior approves the licences. Everyone receives a login on Monday. On Wednesday there is a one-hour session called Introduction to AI, with biscuits. By Friday, half the staff have asked the machine to rewrite an email and the other half have forgotten the password.

Try the same plan at an airline. Give the captain, cabin crew, engineer and baggage handler the same manual and the same afternoon of training. The baggage handler learns about crosswinds. The captain learns how to stack suitcases. Everyone has been trained, yet nobody does their own job better.

Work happens through roles. A bookkeeper reconciles accounts, a salesperson pursues quotes, and an operations manager keeps the promises the salesperson made. Each role needs AI for different tasks, with different context and a different standard for a good answer.

In a small business, one person may hold several roles before lunch. The roles remain distinct. Each has different inputs, outputs, risks and decisions.

The same DataCamp survey found that 23% of leaders said their AI training was not tailored to specific roles. That is the gap RiBAIA addresses.

Role-Based AI Adoption starts with the work and the role accountable for it. The choice of software comes later.

TODAY AND TOMORROW

From the apron to the beach

Two airlines bought the same jet in the same month.

The first airline looks like many companies today. The jet idles on the apron. Staff climb in and out when they feel like it, pressing whichever buttons look friendly. One person teaches the autopilot a clever trick and tells nobody. The finance director sees a fuel bill every month and no ticket sales. When the board asks where the aircraft has been, the answer is a list of people who enjoyed sitting in it.

In a business, the same scene looks less funny. Licences are paid for and logins are issued. A few enthusiasts get good results and keep their methods in private chat histories. Work reaches clients without a clear reviewer. Management reports logins because it cannot report value.

The second airline runs a schedule. Every flight has a destination and date. Every seat on the flight deck has a named person with the right licence. The crew know what the autopilot may do and what the captain must decide. The flight plan is filed, the logbook is kept, someone signs for the aircraft, and the instruments are checked before take-off. When the board asks where the aircraft has been, the answer is a list of landings.

RiBAIA ties a two-week target to an annual goal. Roles define what AI may produce, assist or leave human. Prompts are versioned and outputs owned. After ninety days, delivery and value are compared with a baseline.

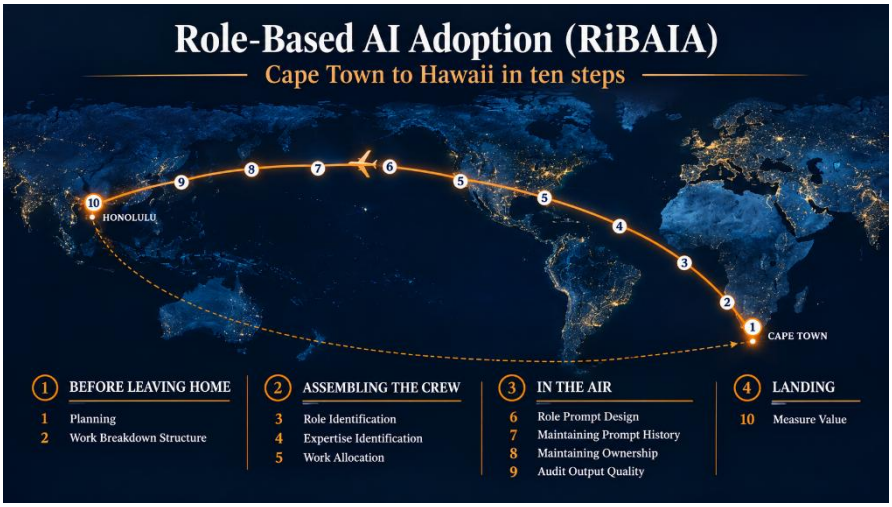
	Today on the apron	After RiBAIA on a schedule
Direction	Start using AI	A two-week target tied to an annual goal
Work	Large, vague requests	Small tasks with one owner each
People	One course for everyone	An adoption plan for each role
Expertise	Act as an accountant	A specific expertise profile
Who does what	Decided by habit	Classified by output and risk
Prompts	One line typed from scratch	A five-part reusable brief
Know-how	Private chat histories	A shared versioned library
Accountability	The AI wrote it	A named owner who reviews and approves
Quality	Errors found by the client	Three checks before use
Proof	Logins counted	Results compared with a baseline

Both airlines own the same jet and employ the same people. The difference is an operating method. The ten steps turn unsupervised experimentation into controlled delivery.

The change takes more than an afternoon. A new pilot flies short routes before long ones. Start the same way: one role, one two-week target and one useful prompt. Then repeat.

THE FRAMEWORK

The route



Role-Based AI Adoption puts AI to work one role at a time. Its ten steps connect business intent to small tasks, assign accountable roles, divide work between people and AI, and test quality and value.

Stage of the trip	Step	On the journey
Before leaving home	1 Planning	Choose the business destination and two-week target
	2 Work Breakdown Structure	Find hidden work and split it into finishable tasks
Assembling the crew	3 Role Identification	Assign every task to a role and named person
	4 Expertise Identification	Describe the expertise each role needs
	5 Work Allocation	Set the boundary between AI and human work
In the air	6 Role Prompt Design	Write a five-part brief for one real output
	7 Maintaining Prompt History	Save and version prompts in a shared logbook
	8 Maintaining Ownership	Keep a named human accountable for every output
	9 Audit Output Quality	Check voice, accuracy and judgment
Landing	10 Measure Value	Compare delivery and business value with the baseline

I developed RiBAIA within EPM-AI, my wider framework for AI-enabled delivery. The named tools in the next ten chapters come from that framework.

RiBAIA uses the controls and software you already have. Each step attaches to a practice such as planning, sign-off or quality review. The summaries mention Jira and Asana, but any capable work-management system will do.

01

Step 1 Planning

Choose the destination and set the date



Ask a travel agent for somewhere warm and you receive a brochure. Ask for Honolulu, arriving on 15 December, and you can plan a route.

Most AI adoption starts with the equivalent of somewhere warm. Leaders say start using AI, so staff use it to produce the same reports and slide decks a little faster. The aircraft is moving, but it is circling the car park.

Step 1 decides where the work must land before anyone touches the controls. RiBAIA uses a simple planning chain: vision, multi-year strategy, annual goal, quarterly objective, two-week target and the tasks needed to achieve it.

AI acts at the bottom of that chain. It can help complete tasks. It cannot rescue a target that does not support a business objective.

Use one test. Pick a person and ask what they have committed to finish today. Then trace that task to the strategy. Where the line breaks, AI risks accelerating work the business does not need.

A small accounting practice wants 40 new monthly bookkeeping clients this year. Its quarterly objective is to cut onboarding time from three weeks to one. Its next two-week target is to create a standard onboarding pack for retail clients. The destination is now specific enough to plan.

STEP 1 AT A GLANCE	
YOU START WITH	Your strategy and your goals for this year.
YOU DO	**1 Name the annual goal.** **2 Set the quarterly objective.** **3 Write one two-week target.** **4 Trace that target to the strategy.**
EPM-AI TOOL	Entinology Planning Hierarchy
FITS WITH	Strategic planning, OKRs and quarterly or sprint planning.
LIVES IN	A goal in Asana, or an epic with a sprint goal in Jira.
YOU END WITH	**One two-week target tied to a business goal.**
NEXT	Step 2 Work Breakdown Structure

02

Step 2 Work Breakdown Structure

Cut the journey into work people can finish



Nobody flies from Cape Town to Hawaii in one leg. The journey may run through Johannesburg and Sydney before Honolulu. Each leg breaks down again into passports, visas, transfers, bags and someone feeding the dog.

A two-week target works the same way. Build the onboarding pack is still too large to assign or prompt well. Hand that whole request to AI and it is likely to return something broad and generic.

Break the target into tasks until each task has one owner and can be finished in a focused work session. RiBAIA uses six hours as a practical ceiling, not a law of nature. Label each task as keyboard time, activity time or wait time. Plans fail when they count desk work but ignore workshops, approvals and delays.

Before estimating, look for missing work. Experienced travellers do this at the front door: passport, wallet, phone, keys. RiBAIA uses the Entinology Diamond for the same purpose.

The Diamond asks whether the target creates work in nine areas: technology, process, customer, channel, data, regulation, organisation, software and location.

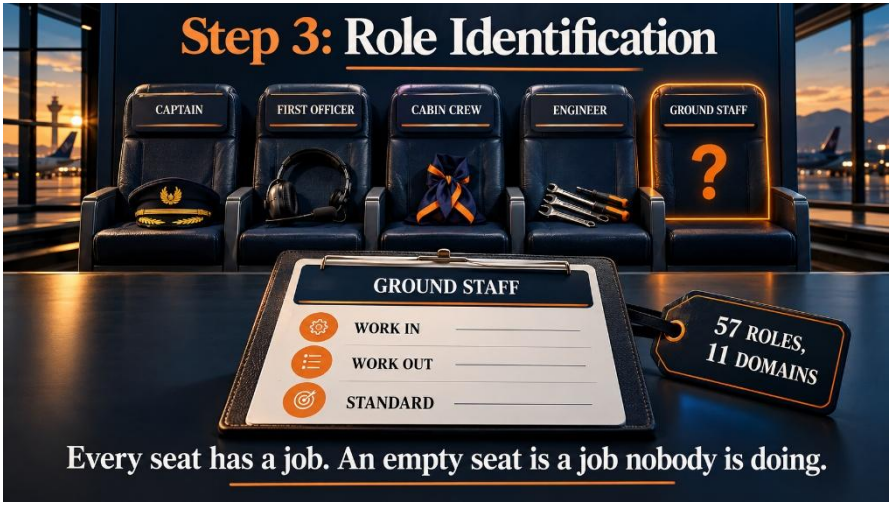
For the onboarding pack, the data question reveals old client records that need cleaning. The regulation question reveals consent wording for personal information. Finding both in week one is planning. Discovering them in week two is delay.

STEP 2 AT A GLANCE	
YOU START WITH	The two-week target from step 1.
YOU DO	**1 Use the nine Diamond areas to find hidden work.** **2 Split work into tasks small enough to finish.** **3 Give each task one owner and a time type.**
EPM-AI TOOL	Entinology Diamond and the three time types
FITS WITH	Backlog planning, sprint planning and impact assessment.
LIVES IN	Tasks and subtasks in Asana, or stories and sub-tasks in Jira, each with an assignee and a time estimate.
YOU END WITH	**A complete task list with owners, time types and estimates.**
NEXT	Step 3 Role Identification

03

Step 3 Role Identification

Name every seat before take-off



A long-haul flight needs a captain, first officer, cabin crew, engineers and ground staff. Leave one seat unassigned and the work attached to it will surface at the worst possible time.

Take the task list from Step 2 and ask which role owns each task. A role is not a person, team or department. It is a defined bundle of inputs, outputs, decisions and standards. Once the role is clear, assign a named person to it.

That definition gives AI useful context. Our team of eight doing various things is vague. A bookkeeper who receives bank statements, produces reconciled accounts and answers to the practice owner gives the system a much narrower job.

The role also answers the question a client, auditor or regulator will ask: who is accountable for this output? Software can have a user. Accountability belongs to a named person acting in a defined role.

EPM-AI maps 57 roles across 11 business areas. A large organisation may assign each role to a different person. In a six-person company, the owner may be captain, navigator and baggage handler. List the roles separately anyway. When using AI, state which seat you occupy.

The accounting firm's onboarding pack involves four roles: practice owner, bookkeeper, client manager and compliance reviewer. Two roles happen to be held by the same person.

STEP 3 AT A GLANCE	
YOU START WITH	The task list from step 2.
YOU DO	**1 Define the role's inputs, outputs, decisions and standard.** **2 Assign one named person to each role for this work.** **3 Mark where one person holds several roles.**
EPM-AI TOOL	The EPM-AI role catalogue: 57 roles in 11 areas
FITS WITH	Job descriptions, the organisation chart and RACI.
LIVES IN	The assignee, plus a Role field or label on every task in Jira or Asana.
YOU END WITH	**Every task linked to a defined role and named person.**
NEXT	Step 4 Expertise Identification

04

Step 4 Expertise Identification

Match the licence and experience to the seat



Imagine boarding after this announcement: Your pilot today has seen several aeroplanes. You would want more detail: which aircraft, how many hours and whether any were flown over an ocean at night.

AI needs specific context too. Act as an accountant invites a generic answer. A senior bookkeeper serving retail clients, familiar with sales tax rules and owners who keep receipts in a shoebox, gives the system a clearer frame.

A detailed expertise description does not make AI an expert. It improves the relevance of the draft and gives the human reviewer a sharper basis for judging it.

For every role, record four parts of its Role Context: discipline, seniority, industry knowledge and the specialised knowledge this task requires.

Then define the work. The full EPM-AI tool is called the Caddie's Ten because, like a good golf caddie, it makes the player think before swinging. In this field guide, capture at least the task's purpose, intended user, required expertise, quality standard, checker and accountable owner.

Those answers take minutes to write and remove much of the ambiguity that causes weak AI output.

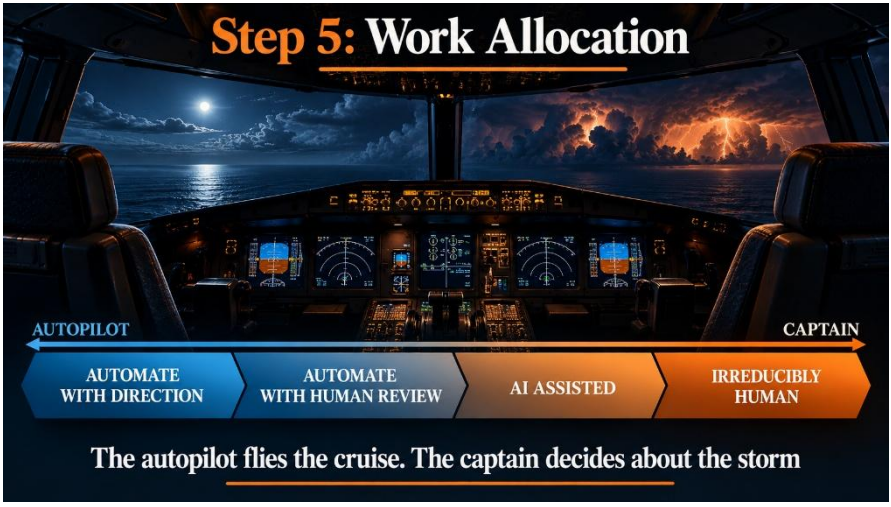
The accounting firm's profile reads: senior bookkeeper; retail clients; sales tax and payroll; proficient in the practice's accounting software; writes to clients in plain language.

STEP 4 AT A GLANCE	
YOU START WITH	The roles from step 3.
YOU DO	**1 Record discipline, seniority, industry and specialist knowledge.** **2 For each task, capture purpose, user, required expertise, quality, checker and owner.**
EPM-AI TOOL	Role Context and the Caddie’s Ten
FITS WITH	Job profiles, the skills matrix and acceptance criteria.
LIVES IN	A role profile page in Confluence, Notion or SharePoint, linked from the tasks.
YOU END WITH	**A concise expertise profile and task context for each role.**
NEXT	Step 5 Work Allocation

05

Step 5 Work Allocation

Decide what the autopilot may fly



On a long ocean crossing, the autopilot does much of the routine flying. The captain still owns the flight, handles take-off and landing, and decides whether to divert around a storm.

Every role needs the same explicit boundary. Without it, people give AI decisions it should not make and keep doing routine production work that AI could handle under supervision.

RiBAIA assigns each recurring output to one of four groups. Classify the output, not the job title.

- **Automate with Direction. AI produces the draft from a human brief. A person checks it before use. Suitable work is routine, reversible and governed by a clear standard.**

- **Automate with Human Review. AI produces a first version, and a qualified professional verifies it against authoritative sources because an error could cause material harm.**
- **AI Assisted. A person leads the work and uses AI for defined parts such as research leads, options or drafting. The person controls the reasoning and conclusion.**
- **Irreducibly Human. A person performs the work. Decisions with serious consequences, legally binding commitments, difficult conversations and ethical judgments stay here.**

For the bookkeeper, a monthly document reminder is Automate with Direction. A bank reconciliation is Automate with Human Review. Cash-flow options are AI Assisted. Telling an owner that the business is insolvent remains human, in person, with the door closed.

Update the Step 2 estimates and define the review for every AI-assisted output. Because classification is task-based, the method remains independent of any supplier.

STEP 5 AT A GLANCE	
YOU START WITH	The roles, profiles and tasks from steps 2 to 4.
YOU DO	**1 List each role's recurring outputs.** **2 Classify every output into one of four groups.** **3 Define the review level and update time estimates.**
EPM-AI TOOL	The Deliverable Classification
FITS WITH	Delegation of authority, decision rights and your AI use policy.
LIVES IN	An AI Allocation field on each task in Jira or Asana, with the four groups as options.
YOU END WITH	**A one-page allocation and review table for each role.**
NEXT	Step 6 Role Prompt Design

06

Step 6 Role Prompt Design

File a flight plan the machine can follow



An aircraft follows the route entered into its systems. A wrong waypoint can send an expensive machine smoothly and confidently in the wrong direction.

A prompt is a working brief. Write a welcome pack for new clients is not a brief, yet people are surprised when it produces a generic result.

Build the prompt from five elements in this order.

1. **Role Context. State the role and relevant expertise from Step 4.**
2. **Diamond Activation. Name the Diamond areas from Step 2 that matter to this task, such as data or regulation.**
3. **Caddie's Ten Application. Include the answers that matter most for this task: purpose, user, expertise, quality, checking and ownership.**

4. **Task Specification. State what to produce, for whom, in what format and for which decision or action.**
5. **Quality Standard. Give the test the output must pass before it comes back to you.**

End with a clarification rule: identify missing information and ask questions before drafting. AI will often fill gaps unless you require it to stop.

The bookkeeper's shortened prompt reads like this:

> *You are a senior bookkeeper with fifteen years of experience serving retail clients. This task involves client data and personal-information rules. Draft a two-page welcome pack that a new retail client can read in ten minutes. Use plain language. Include every document on our onboarding checklist and the date by which we need it. The practice owner will verify the draft against the checklist and current policy. Before drafting, list any missing information and ask me to supply it.*

Writing a precise brief takes a few minutes. Repairing an ambiguous one can take far longer.

STEP 6 AT A GLANCE	
YOU START WITH	The allocation table and the expertise profile.
YOU DO	**1 Choose one output AI may produce or assist with.** **2 Build the brief from the five elements.** **3 Require questions when important information is missing.**
EPM-AI TOOL	The five elements of an EPM-AI prompt
FITS WITH	Briefs, templates and standard operating procedures.
LIVES IN	The prompt saved in the task description or as a template in Jira or Asana.
YOU END WITH	**One complete prompt for one real output.**
NEXT	Step 7 Maintaining Prompt History

07

Step 7 Maintaining Prompt History

Keep the logbook



Every aircraft carries a logbook. Crews record faults, conditions and fuel use. The next crew reads the record before starting the engines and avoids relearning the same lesson.

Many businesses use AI without a logbook. A prompt that produces excellent work stays in one employee's chat history. When that person leaves, takes holiday or clears the conversation, the next user starts again from write a welcome pack for new clients.

I call this AI debt: useful AI know-how that has not been documented, versioned or assigned to an owner. It becomes visible when the business must reproduce an output or explain how it was made.

The remedy is a shared prompt library. Save only prompts worth reusing, and record seven fields for each one.

Field	What goes in
Prompt	The full text
Version	1, 2, 3
Date	When this version was saved
Owner	One named person
Role	The seat the prompt was written for
AI tool	Which tool produced the good result
What changed and why	The lesson from the last version

The final field, what changed and why, turns storage into learning.

Start a fresh conversation when a long AI session becomes hard to follow. Carry forward a clean summary, the approved source material and the current prompt. A longer chat is not automatically a better record.

STEP 7 AT A GLANCE	
YOU START WITH	The finished prompt from step 6.
YOU DO	**1 Save the prompt in a shared library.** **2 Record all seven fields.** **3 Start a fresh session with a clean summary when a conversation becomes unwieldy.**
EPM-AI TOOL	The shared prompt library
FITS WITH	Document control, version control and knowledge management.
LIVES IN	A prompt library page in Confluence, Notion or SharePoint, with page history switched on.
YOU END WITH	**A versioned logbook with an owner and change history.**
NEXT	Step 8 Maintaining Ownership

08

Step 8 Maintaining Ownership

Keep a captain above the loop



Before departure, the captain accepts responsibility for the aircraft. The autopilot may operate the controls, but it does not inherit accountability.

AI makes this easy to forget. A draft arrives quickly and sounds polished. The approve button is close. If the output goes out under your name, you own the consequence; the machine will not take the client's call.

EPM-AI calls the correct position Above the Loop. The human sets direction, supplies context, reviews the output, approves its use and intervenes when the risk or evidence demands it.

Ownership is not a cosmetic edit. The owner must read the whole output, verify it to the required standard, correct errors and record approval. An unchanged output may be acceptable; an unread output is not.

Every task from Step 2 already has a named owner. Keep that name attached when AI does part of the work. Give every library entry an owner too, and transfer ownership explicitly when someone leaves.

There is also a human reason for this discipline. In my experience, people are more likely to exercise judgment when they have accepted a clear commitment than when they have merely received a general instruction.

At the accounting firm, the bookkeeper owns the accuracy of the welcome pack. The practice owner owns the decision to approve it for client use.

STEP 8 AT A GLANCE	
YOU START WITH	The logbook and the task list.
YOU DO	**1 Name the owner of every AI-assisted output.** **2 Read and verify the whole output.** **3 Record approval and transfer ownership explicitly when needed.**
EPM-AI TOOL	Above the Loop
FITS WITH	Sign-off, approvals and the audit trail.
LIVES IN	The assignee and approver on the task, with an In Review status in the Jira or Asana workflow.
YOU END WITH	**A named human accountable for each output and prompt.**
NEXT	Step 9 Audit Output Quality

09

Step 9 Audit Output Quality

Check the instruments before trusting them



Pilots use instruments and still cross-check them. Before flight, someone walks around the aircraft. In the air, an odd reading is compared with another source before the crew acts.

AI needs the same discipline because confidence is not evidence. A correct statement, an invented regulation and a misspelled client name can arrive in the same fluent voice.

Run three checks before using any output. RiBAIA calls this the Three-Check Output Filter.

- **Voice. Does this sound like the person or organisation sending it, without inflated claims or generic filler?**

- **Accuracy. Has a person verified every material fact, figure, name and citation against an authoritative source?**
- **Judgment. Is the conclusion appropriate for the audience, risk and decision, and would the owner defend it face to face?**

For higher-stakes work, add independent challenge. Ask the AI to test the output against the stated quality standard, then use a second tool or reviewer to look for different weaknesses. Agreement between AI tools is not proof. Verify material claims against primary sources or a qualified human.

Keep an error log. Record what the check caught, how serious it was and which type of document contained it. After a month, you will know where controls need to be stronger.

Also count how many attempts each prompt needs before the output is usable. A falling number suggests improvement, but only if the quality standard stays constant.

STEP 9 AT A GLANCE	
YOU START WITH	An output with an owner, ready to send.
YOU DO	**1 Check voice, accuracy and judgment.** **2 Add independent challenge for higher-stakes work.** **3 Verify material claims against authoritative sources.** **4 Log errors and attempts.**
EPM-AI TOOL	The Three-Check Output Filter
FITS WITH	Quality assurance, peer review and the definition of done.
LIVES IN	A three-item checklist on the task, and an AI Error label you filter on in Jira or Asana.
YOU END WITH	**A verified output and evidence about recurring error patterns.**
NEXT	Step 10 Measure Value

10

Step 10 Measure Value

Prove you landed somewhere worth going



	BEFORE	AFTER 90 DAYS
WORK FINISHED VS PROMISED	32%	87%
BLOCKER TIME	28%	5%
WAIT TIME	40%	8%
HOURS RELEASED	0	1,200
COST AVOIDED	$0	$180,000
INCOME GAINED	$0	$450,000

The wheels touch down and the cabin applauds. The captain asks less glamorous questions: Is this Honolulu? Are we on time? How much fuel remains? Did every bag arrive?

Speed alone proves little. Fly fast in the wrong direction and you arrive early in Perth.

AI initiatives are often judged by hours saved. That measure is incomplete because it does not show what happened to the released time or whether the work improved.

Record the baseline before the pilot starts and compare it with the same measures after ninety days. RiBAIA tracks delivery measures such as these:

- **Commitment reliability.** Work completed compared with work promised in each two-week period.
- **Blocker time.** Time from a blocker being raised to being cleared.
- **Wait share.** Waiting time as a proportion of the task's total elapsed time.

Add business value: capacity released and how it was used, cost avoided, income accelerated or gained, and quality or risk changes. Do not count theoretical savings as cash unless spending actually falls.

In one earlier implementation of the precursor delivery methods, teams increased completed commitments from about 20% to more than 87% and held about 80% fewer meetings with the same people. These are results from one implementation, not a forecast. Outcomes depend on the organisation and execution.

For the accounting firm, the baseline is concrete: three weeks to onboard a client, eleven emails per onboarding and four new clients per month. After ninety days, the same three measures show whether the journey improved delivery. Financial measures show whether it was worth the fuel.

The evidence from Step 10 becomes the input to the next planning cycle. RiBAIA is a loop, not a once-off rollout.

STEP 10 AT A GLANCE	
YOU START WITH	Checked work going out, and today's numbers.
YOU DO	**1 Record the baseline before the pilot.** **2 Compare the same measures after ninety days.** **3 Measure capacity, cost, income, quality and risk.** **4 Use the evidence in the next plan.**
EPM-AI TOOL	The EPM-AI delivery measures
FITS WITH	KPI reporting, benefits tracking and retrospectives.
LIVES IN	A dashboard in Jira or Asana: work completed against committed, time blocked and cycle time.
YOU END WITH	**Before-and-after evidence that feeds the next planning cycle.**
NEXT	Back to Step 1 Planning

THE PROCESS

The whole flight on one page

Each step produces an output that becomes the next step's input. Read down the final column to see the full adoption process from business goal to evidence.

	Step	You do	You end with
01	**Planning**	Set a two-week target tied to strategy	**Two-week target**
02	**Work Breakdown Structure**	Find hidden work and split it into small tasks	**Task list**
03	**Role Identification**	Link every task to a role and person	**Roles and owners**
04	**Expertise Identification**	Describe the expertise and task context	**Expertise profiles**
05	**Work Allocation**	Set the boundary between AI and people	**Allocation tables**
06	**Role Prompt Design**	Build a five-part brief	**Complete prompt**
07	**Maintaining Prompt History**	Save and version prompts	**Shared logbook**
08	**Maintaining Ownership**	Review, verify and approve	**Named accountability**
09	**Audit Output Quality**	Check voice, accuracy and judgment	**Verified output**
10	**Measure Value**	Compare results with the baseline	**Evidence for the next cycle**

ARRIVAL

On the beach

The crew reach Waikiki. The captain still has the logbook open on a sun lounger, to the mild irritation of everyone nearby. The next cycle starts in ninety days, and it should use less fuel for a better result.

You now know the route: choose useful work, break it down, assign the right roles, define expertise, divide work between people and AI, design and preserve prompts, keep human ownership, check the output and measure value. Start with one role and one target. Earn the right to scale.

Sources for opening statistics: DataCamp, The AI Skills Gap in 2026, 12 March 2026, datacamp.com; Boston Consulting Group, BCG AI Radar 2026, 15 January 2026, bcg.com.

THREE DAY ROLE BASED AI ADOPTION WORKSHOP

Build your first RiBAIA pilot in three days

Bring one real business target and the people who own it. I will lead your team through all ten steps using your own work.

YOU WILL LEAVE WITH

- A two-week target tied to strategy
- A task map, role profiles and AI work-allocation table
- Your first tested prompt saved in a shared logbook
- A 90-day baseline for quality, capacity and value

By the final session, your first pilot is ready to run. Maximum 15 seats. I lead every session.

BOOK YOUR PLACE

entinology.com | WhatsApp +27 83 499 5112

Current dates, format, venue and fees are available online.

THE AUTHOR

About the author



Roelof Vermeulen

Founder of Entinology. Creator of EPM-AI and RiBAIA.

Roelof Vermeulen is the founder of Entinology and the creator of EPM-AI and Role-Based AI Adoption. He developed RiBAIA to give organisations a practical way to assign AI work, human judgment and accountability one role at a time.

He has nearly 30 years of delivery experience across banking, telecoms, insurance, retail, hospitality, software and higher education on several continents. He has written five books on AI and project delivery and speaks to executive and project audiences around the world.

Roelof designs and leads every engagement, supported by the Entinology team.

First published October 2026. Copyright 2026 Roelof Vermeulen. All rights reserved. Published by Entinology (Pty) Ltd.