AXIA LEARNING / AIPM

Diploma in AI Product Management

Five months. Practical learning. A clear next step.



20 weeks | **220 hours** | **Live classes**

Developing Talent. Building Intelligent Solutions.

Your pathway



Focus and scope. Product discovery, strategy, delivery and measurement for AI-enabled products. The focus is deciding what should be built, proving value and managing risks, rather than training models.

Entry requirements. Basic spreadsheets, communication and business literacy. No coding requirement; a short diagnostic checks numeracy and problem framing.

Pre-course preparation. Optional 8–12-hour preparation in spreadsheet formulas, product terminology and writing structured business documents.

Primary tools. Core: documents, spreadsheets, Figma viewer, a backlog tool and one AI assistant. Use sample analytics data and supplied model-evaluation outputs; no production integrations required.

Career direction. Portfolio preparation for product internships, product operations and associate/junior product applications; senior product ownership requires further experience.

What you'll be able to do

LO1. Discover a worthwhile user problem and assess whether AI is appropriate.

LO2. Write an AI product strategy, business case, scoped PRD and measurable acceptance criteria.

LO3. Compare build, buy and non-AI options using cost, latency, quality and operational constraints.

LO4. Design a data and evaluation plan, including failure cases and human review.

LO5. Manage a delivery backlog, cross-functional trade-offs and a staged rollout.

LO6. Interpret product metrics and defend a launch, change or stop decision.

Five-month overview

Month	Focus	Main assessment
1	Product discovery and foundations	Opportunity assessment with research evidence and AI/non-AI comparison.
2	Strategy, commercial reasoning and requirements	Product strategy, PRD, prioritised roadmap and unit-economics scenarios.
3	AI delivery, data and evaluation	AI delivery plan, data-readiness assessment and evaluation design.
4	Measurement, launch and operations	A separate product-launch/post-launch case using supplied data and a decision memo.
5	Capstone and product defence	Final capstone and individual practical

MONTH 1 / 05

Product discovery and foundations



1. Product management and AI fundamentals

Skills & concepts

Distinguish product outcomes from outputs; understand AI limitations and common product patterns.

Practical evidence

Critique an AI feature and propose a measurable customer outcome.

2. Customer discovery

Skills & concepts

Plan interviews, identify user segments and distinguish buyers, users and stakeholders.

Practical evidence

Conduct at least three discovery interviews and document supporting evidence.

3. Problem validation

Skills & concepts

Map journeys, assess workarounds and estimate problem frequency and importance.

Practical evidence

Build an opportunity assessment with assumptions and evidence gaps.

4. Discovery decision

Skills & concepts

Compare rule-based, manual and AI approaches; define an initial success metric.

Practical evidence

Submit Month 1 opportunity brief and a reasoned go/no-go recommendation.

MONTH 2 / 05

Strategy, commercial reasoning and requirements



5. Product strategy

Skills & concepts

Define positioning, target segment, value proposition and strategic trade-offs.

Practical evidence

Write a one-page strategy with explicit exclusions and competitive alternatives.

6. Business models and unit economics

Skills & concepts

Model willingness-to-pay assumptions, cost per successful task and support/review costs.

Practical evidence

Build a scenario model with adoption, model-use and human-review sensitivity.

7. Prioritisation and roadmap

Skills & concepts

Rank opportunities, compare confidence levels and sequence discovery and delivery.

Practical evidence

Produce a prioritised roadmap linked to outcomes rather than feature count.

8. PRDs and acceptance criteria

Skills & concepts

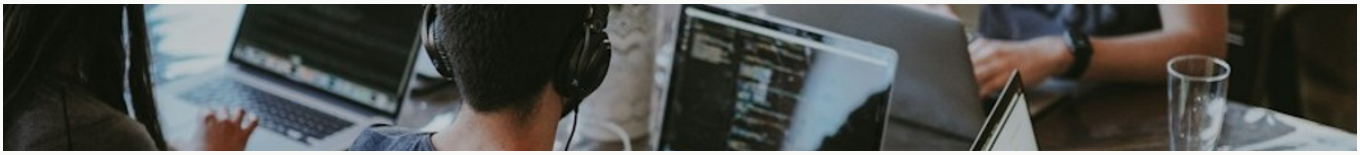
Specify user flows, edge cases, instrumentation and release boundaries.

Practical evidence

Submit Month 2 PRD and business case with testable acceptance criteria.

MONTH 3 / 05

AI delivery, data and evaluation



9. AI architecture for product managers

Skills & concepts

Understand model APIs, retrieval, tools, latency and build/buy trade-offs at product level.

Practical evidence

Produce an architecture-options brief with engineering questions and constraints.

10. Data readiness and governance

Skills & concepts

Identify data sources, rights, labelling needs, quality gaps and access controls.

Practical evidence

Create a data-readiness checklist, ownership map and risk register.

11. Evaluation and quality gates

Skills & concepts

Define representative test cases, user outcomes, factuality and task completion criteria.

Practical evidence

Build a 30-case evaluation plan with baseline, slices and failure handling.

12. Cross-functional delivery

Skills & concepts

Write stories, manage dependencies and negotiate scope with design and engineering.

Practical evidence

Submit Month 3 delivery backlog, evaluation gates and release definition.

MONTH 4 / 05

Measurement, launch and operations



13. Product analytics and experiments

Skills & concepts

Define events, funnels, cohorts, hypotheses, guardrails and uncertainty.

Practical evidence

Analyse supplied product data and draft an experiment with a stopping rule.

14. Responsible rollout

Skills & concepts

Set human oversight, acceptable use, escalation, consent and staged release criteria.

Practical evidence

Produce a pilot plan with incident ownership and rollback triggers.

15. Commercial launch and adoption

Skills & concepts

Plan onboarding, positioning, customer education, sales enablement and support.

Practical evidence

Deliver a launch checklist and an adoption experiment.

16. Lifecycle operations

Skills & concepts

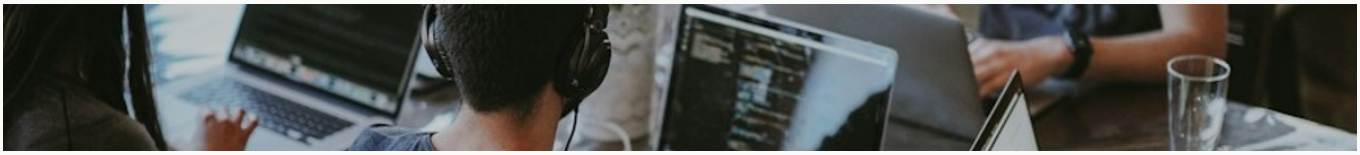
Review model/vendor changes, support signals, cost drift and deprecation decisions.

Practical evidence

Submit Month 4 post-launch case analysis; approve a separate capstone scope.

MONTH 5 / 05

Capstone and product defence



17. Capstone opportunity and strategy

Skills & concepts

Validate a selected problem and establish scope, business assumptions and success measures.

Practical evidence

Submit an evidence-backed product opportunity and strategy pack.

18. Capstone specification

Skills & concepts

Integrate requirements, prototype, evaluation plan and realistic delivery dependencies.

Practical evidence

Present PRD, prototype walkthrough and prioritised backlog.

19. Capstone launch simulation

Skills & concepts

Use an evaluation dataset and cost model to respond to a launch-readiness scenario.

Practical evidence

Submit a release recommendation, risk response and rollout plan.

20. Product defence and career preparation

Skills & concepts

Explain trade-offs, answer stakeholder objections and interpret a new data scenario.

Practical evidence

Deliver a product portfolio and individual practical defence.

Your assessment checkpoints

Checkpoint	Assessed project	Outcome coverage
Week 4	Opportunity assessment with research evidence and AI/non-AI comparison.	LO1
Week 8	Product strategy, PRD, prioritised roadmap and unit-economics scenarios.	LO2, LO3
Week 12	AI delivery plan, data-readiness assessment and evaluation design.	LO4, LO5
Week 16	A separate product-launch/post-launch case using supplied data and a decision memo.	LO6

The capstone and final individual practical together verify all six outcomes. Instructors use the weekly evidence to identify gaps early; the mapping above does not limit remediation to those weeks.

Your final capstone

Own a proposed AI product from discovery to launch readiness. Submit a product strategy, research evidence, PRD, prototype, roadmap, cost model, evaluation plan and staged rollout recommendation.

Capstone acceptance criteria

Support the problem with at least five capstone user/stakeholder interviews, or an approved access-limited research alternative clearly labelled. Include at least 30 evaluation cases with expected outcomes, critical failures and launch gates. Compare at least two implementation options and one non-AI alternative. Provide cost-per-successful-task scenarios, guardrail metrics and rollback ownership. Defend decisions in an individual review; a launch plan alone is not evidence of a successful launch.

Evidence pack. Include a short brief, source/permissions record, versions, practical artefact, testing or user-research evidence, limitations, handover instructions and a record of material AI assistance. Include a 5–8-minute demonstration and an individual 15-minute defence inside the final week's facilitated schedule.

How the programme works

Programme model and delivery

Programme structure. Six separate pathways, each lasting 20 teaching weeks, organised into five four-week learning blocks. A student enrols in one diploma. The five-month calendar should be confirmed against holidays and local delivery dates. This is a proposed institutional curriculum blueprint, not an accreditation decision.

Audience and admission. Mixed cohort: beginners and career changers can enter the non-coding tracks after an orientation. Technical tracks use practical entry checks and pre-course bridging where required. AI Engineering requires Python/API/Git competence before Week 1. Assess applicants by demonstrated skills rather than degree title alone.

Weekly workload: 11 hours. 3 hours of interactive teaching; 2 hours of supervised practical work; 1 hour of critique, review or mentoring; 5 hours of independent reading, building and assessment preparation. Total: 220 learning hours per diploma, including 120 facilitated hours and 100 independent hours. Projects and assessment preparation are included, not added on top.

How each week runs. Begin with a concrete work problem and a short retrieval exercise; teach the relevant concepts; demonstrate a worked example; complete a supervised lab; finish an individual deliverable and receive feedback. Use a different dataset or brief for independent work so students cannot pass by copying the demonstration.

Shared foundation without losing track time. Week 1 in each course includes AI limitations, data privacy, responsible tool use and evidence checking through a track-specific exercise. Communication, version history, documentation and critical thinking recur throughout. Do not add a separate generic month that displaces the core discipline.

Progression gates. Week 4: demonstrate foundational work independently. Week 8: complete a usable core project. Week 12: show role-specific AI application and evaluation. Week 16: demonstrate integrated work and obtain approval for a separate capstone. Weeks 17–20: deliver, test and defend the capstone. Learners who miss a gate receive targeted remediation and a reassessment plan.

AI use in assessment. Allow AI in normal labs if students disclose material use and verify outputs. Keep a brief record of prompts, generated material and corrections. Do not accept fabricated interviews, invented datasets, unrun tests or unsupported impact claims. The final practical uses a fresh task, allows reference documentation and excludes AI-generated solutions; the oral defence verifies individual understanding.

Data and lab practice. Use public, licensed, synthetic or appropriately consented data. Do not place confidential organisational or personal data in external AI tools. Cybersecurity practice takes place only in written-scope isolated labs; automation uses sandbox accounts and approval before external side effects. Students keep a source and permissions register.

Capstone format. Prefer individual work for a five-month part-time diploma. If a team of up to three is approved, each learner keeps a contribution log and submits an individually assessed component and defence. A cross-track showcase is optional; a

learner's grade must not depend on another track delivering its component.

Portfolio and employability. Career preparation is embedded in project write-ups and mentor review. Each student finishes with four graded monthly project artefacts, one capstone case study, a concise CV/profile and a rehearsed demonstration. Describe evidence of competence honestly; the curriculum does not promise employment or advanced professional expertise.

Assessment and progression

Assessment	Weight	Evidence
Weekly practical work, Weeks 1–16	20%	Sixteen individual practical submissions; mark correctness, reasoning and verification. Each contributes 1.25%.
Four monthly projects, Weeks 4, 8, 12 and 16	30%	Each contributes 7.5%. Grade the integrated project against its outcome map; do not award the same points twice for a weekly lab.
Capstone, Weeks 17–20	35%	A separate brief with implementation/prototype, evaluation evidence, documentation and individual defence. Reusing earlier techniques is allowed; copied project submissions are not a new capstone.
Final individual practical, Week 20	15%	An unseen role-specific task, normally 90 minutes, followed by a short explanation of decisions. Scheduled inside the weekly facilitated hours.

Proposed capstone rubric (100 points within the capstone component):

Criterion	Points
Problem understanding and scope	10
Discipline-specific execution	35
Testing, evaluation and evidence	20
User/business relevance	10
Responsible practice and risk controls	10
Documentation and reproducibility/handoff	10
Presentation and defence	5

The discipline-specific acceptance criteria in each course are additional gates. A polished presentation cannot compensate for missing core functionality, unsupported findings or lack of individual understanding.

Delivery & learning support

Delivery preparation and quality assurance

Before enrolment. Publish entry requirements, assessment rules, actual tool/API charges and device requirements. Run the diagnostic and bridge. Check access to every compulsory tool on the learner's device. No learner should discover a mandatory paid subscription halfway through the course.

Teaching assets required. For every week prepare one instructor plan, a worked example, a supervised lab brief, an independent task, starter files or data, an answer/review guide and a marking checklist. For each project prepare a brief, rubric, sample acceptable evidence and an individual-defence prompt.

Instructor and mentor support. Use a named practitioner lead per track and approximately one practical mentor per 20–25 learners; reduce the ratio for engineering or complex lab sessions. Offer a recorded session or equivalent guided task for

learners with connectivity constraints.

Device and access baseline. Recommend a laptop with at least 8 GB RAM; provide browser labs or a higher-memory option for virtualisation-intensive work. Cybersecurity learners need isolated lab access; Power BI learners need compatible Windows access or a preselected cohort-wide alternative. AI Engineering does not require students to purchase a GPU.

Tool-cost control. Choose one main tool per capability. Provide sandbox datasets, shared lab infrastructure, bounded API quotas and a mock/offline fallback where feasible. Approve spending limits before live integrations. Verify current pricing and licences before each cohort rather than printing unverified prices in the syllabus.

Feedback and quality assurance. Return weekly feedback within seven days. Moderate a sample of marked work across instructors, verify final individual competence and gather feedback at Weeks 4, 8, 12, 16 and 20. Update tool-specific instructions before the next cohort while keeping learning outcomes stable.

Completion rules proposed for adoption. Require at least 60% overall, 60% in the capstone and 60% in the final practical, plus satisfactory evidence for every core learning outcome. Offer targeted resubmission for unmet outcomes. Require participation in 80% of facilitated activities or approved equivalent tasks; publish any local policy changes before enrolment.

Award and recognition. The provider must approve the final award title, hours and assessment policy. These documents do not establish equivalence to a regulated qualification or

endorsement by the reference organisations. Any accreditation or external-certification claims require a separate, verified approval process.

Reference guidance and curriculum maintenance

These are reference resources used to inform the proposed scope, not course endorsements or a claim of formal certification alignment. Reviewed 17 November 2026. Tool versions, access and pricing must be rechecked before delivery. Human–AI interaction design. Microsoft Research: Guidelines for Human-AI Interaction — https://www.microsoft.com/en-us/research/project/guidelines-for-human-ai-interaction/ AI product risk and evaluation. NIST AI Risk Management Framework and GenAI profile resources — https://www.nist.gov/itl/ai-risk-management-framework Analytics data preparation. Microsoft Learn: Prepare data for analysis with Power BI — https://learn.microsoft.com/en-us/training/paths/prepare-data-power-bi/ Cybersecurity work and skills. NIST NICE Framework Resource Center — https://www.nist.gov/itl/applied-cybersecurity/nice/nice-framework-resource-center AI automation components and evaluation. n8n: Integrate AI — https://docs.n8n.io/build/integrate-ai.md LLM foundations and technical prerequisites. Hugging Face LLM Course — https://huggingface.co/learn/llm-course/chapter1/1 Predictive-model evaluation. scikit-learn: Model selection and evaluation — https://scikit-learn.org/stable/model_selection.html AI application security. OWASP GenAI/LLM application security project — https://owasp.org/projects/top-10-for-large-language-model-applications Human–AI interaction guidance informs expectations, correction and recovery in Design. NIST AI risk guidance informs product evaluation and governance. NICE informs task-based cybersecurity planning. Microsoft data-preparation guidance supports analytics sequencing. n8n and Hugging Face documentation inform automation and engineering capability boundaries. scikit-learn and OWASP inform evaluation and security practice.

Questions? Contact info@axia.africa.