ZEROTOAIMASTERY.COM / ORIGINAL RESEARCH

AI PRODUCT MANAGER HIRING REPORT 2026

What 2,400+ job postings reveal about skills, salaries, degrees and experience

2,400+	1,560	53	30
UNIQUE POSTINGS	COMPANIES	COUNTRIES	CITIES

RESEARCH BY

Abhishek Ashtekar

ZeroToAIMastery.com

DATA SNAPSHOT

12 September 2026

Contents

Twenty-six sections in five parts, built on 2,448 unique postings. Every figure is computed from that dataset; external sources are cited separately and never mixed into a finding.

PART I

The Market

01 Executive summary
02 The AI Product Manager market in 2026
03 What an AI Product Manager actually is

PART II

The Requirements

04 The experience question
05 Do you need a degree?
06 The AI Product Manager skill stack
07 The most requested skills
08 The modern AI stack
09 How technical do AI PMs need to be?

PART III

The Money

10 Salary report
11 The $200K+ AI Product Manager
12 Top industries and employers
13 Remote, hybrid or office
14 Geographic report
15 What high-paying jobs ask for

PART IV

The Signal

16 The skills gaining momentum
17 AI PM role archetypes
18 What employers really want

PART V

The Playbook

19 The job-ready AI PM stack
20 Roadmaps for aspiring AI PMs
21 Portfolio recommendations
22 What to learn before 2027
23 Key takeaways
24 Methodology
25 Data limitations and coverage
26 Appendix

PART I

The Market

How large AI product management has become, where the roles sit, and what the job actually is once you read past the title.

01 Executive summary

02 The market in 2026

03 What an AI Product Manager actually is

SECTION 01

01 Executive summary

What 2,448 unique AI Product Manager job postings say about the skills, pay, credentials and experience employers are actually buying in 2026.

1,560 hiring companies across 53 countries

67% of postings state an experience requirement; 37% disclose pay

Tier A + Tier B postings only; AI-adjacent roles are reported separately

AI product management is no longer an emerging job title. It is a staffed, budgeted, globally distributed function with a recognisable skill profile — and that profile is far more product-led than the discourse around it suggests.

The findings below come from 2,448 unique postings, de-duplicated across employer career sites, applicant-tracking systems and job platforms. Where a number describes the market, it is measured from this dataset. Where a number is contested or partially observed — pay in particular — the coverage rate is stated alongside it.

The ten findings that matter

01 **Product craft still decides who gets hired.** Artificial Intelligence appears in 99% of postings, but the requirement that separates candidates is the product core: Product Strategy (53%), Roadmapping (76%) and Stakeholder Management (64%). Employers are buying product managers who are fluent in AI, not AI specialists who can be taught product.

02 **You do not need to be an engineer.** 84% of postings name no programming language at all. Python appears in 9% and SQL in 11% — meaningful, but far from the universal gate the internet assumes.

03 **Degrees are rarely a hard gate.** Only 22% of postings state a degree as a firm requirement. 53% never mention one, and 11% explicitly allow equivalent experience. An MBA is named in 8% of postings.

04 **This is a mid-to-senior market.** The median posting asks for 5 years of experience. Only 8% of postings that state a figure accept two years or fewer — the single hardest structural fact for career changers to plan around.

05 **Posted pay is high and widely dispersed.** The global median posted midpoint is $193K USD-normalised, with a 90th percentile of $265K. In the United States alone the median is $198K. Coverage: 37% of postings disclose a range.

06 **The frontier stack is a differentiator, not a baseline.** RAG is named in 19.0% of postings, agents in 51.2%, evaluations in 19.6% and fine-tuning in 7.1%. These are minority requirements — which is exactly why they move candidates up shortlists.

07 **Demand is concentrated but not captured.** United States accounts for 52% of postings, yet the top ten employers together account for only 8% of the sample and 76% of hiring companies posted exactly one role.

08 **Hybrid and remote split the market.** 25% of postings are hybrid, 13% fully remote and 15% on-site, with 48% stating no working model at all.

09 **The role has split into distinct archetypes.** Agentic AI PM is the largest at 30% of postings, but seven recognisable archetypes now exist, each with a different skill profile and a different entry path.

10 **The premium segment is a scope story, not a tooling story.** Roles at $200K+ (44% of pay-disclosing postings) ask for 7 years of experience versus 5 years elsewhere, and skew toward platform, infrastructure and organisational leadership rather than any single named technology.

The numbers, at a glance

2,448
unique postings analysed
Tier A + Tier B

1,560
hiring companies
after name normalisation

53
countries represented
employer-stated location

99%
name Artificial Intelligence
most common requirement

5 yrs
median experience asked
67% of postings state one

22%
state a degree as required
53% never mention one

$193K
median posted salary
USD-normalised midpoint

$265K
90th percentile
n=904 disclosing

38%
remote or hybrid
of postings stating a model

Employers are not hiring researchers. They are hiring product managers who are not afraid of models.

WHAT THIS MEANS

The gap between what candidates prepare for and what employers ask for is the most actionable finding in this report. Candidates over-index on model internals and named frameworks. Employers over-index on shipping, judgement under uncertainty, evaluation discipline and the ability to make a probabilistic system safe enough to sell. Section 19 converts this into a prioritised learning stack.

SECTION 02

02 The AI Product Manager market in 2026

How large the market is, where the roles sit, which industries are buying, and who is posting.

United States holds 52% of postings

Enterprise Software & SaaS is the largest hiring sector at 18%

76% of employers posted exactly one AI PM role

The AI Product Manager market in 2026 is broad rather than deep: thousands of employers each hiring one or two people, rather than a handful of laboratories hiring hundreds.

Across 2,448 unique postings this research identified 1,560 distinct hiring organisations in 53 countries. The distribution is the important part. The ten largest employers account for 8% of all postings — meaning roughly 92% of demand sits outside the companies candidates typically target. 76% of hiring companies posted a single AI PM role in the period.

That shape has a direct consequence for job seekers. A search strategy built around frontier AI labs and the largest technology employers addresses a small slice of the market. The larger opportunity is in companies that are adding their first or second AI product manager — where the competition is thinner and the brief is broader.

Where AI Product Manager roles are being posted

Unique 2026 postings by employer-stated country of the role

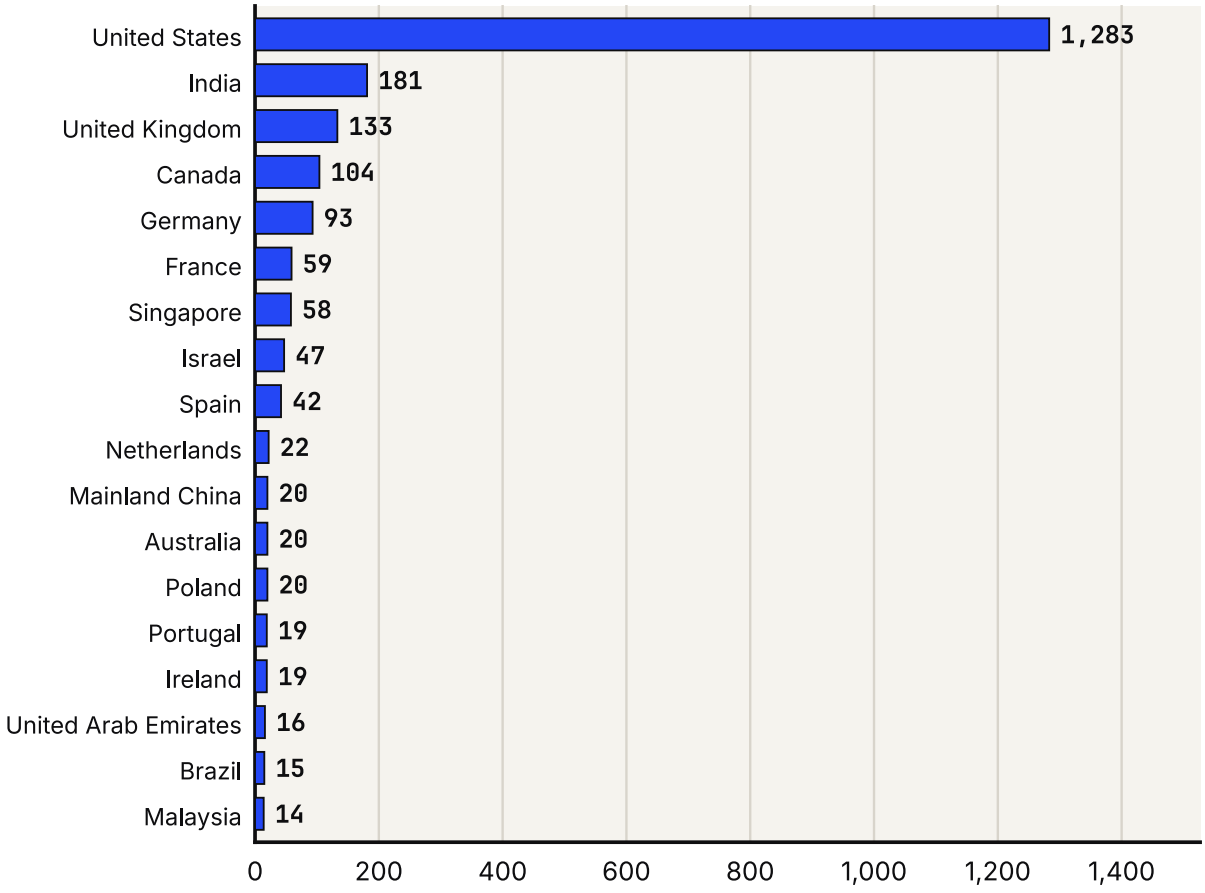

52% of all AI PM demand sits in United States; the next 4 markets together add 21%.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Countries with fewer than 3 postings are omitted.

Figure 1 — Unique 2026 AI Product Manager postings by country.

AI Product Manager demand by world region

Share of unique 2026 postings

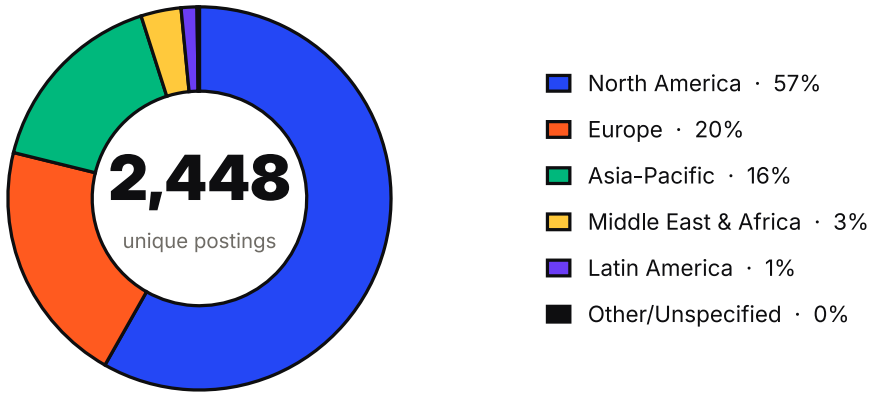

57% of postings originate in North America.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 2 — Regional distribution of AI Product Manager demand.

Where the roles are

United States dominates with 52% of postings (1,283 roles). The next tier — India, United Kingdom, Canada and Germany — adds 21% between them. Beyond that the market is a long tail: 25 countries recorded at least ten unique postings.

COUNTRY	POSTINGS	SHARE	REGION
United States	1,283	52.4%	North America
India	181	7.4%	
United Kingdom	133	5.4%	
Canada	104	4.2%	
Germany	93	3.8%	
France	59	2.4%	
Singapore	58	2.4%	
Israel	47	1.9%	
Spain	42	1.7%	
Netherlands	22	0.9%	
Mainland China	20	0.8%	
Australia	20	0.8%	
Poland	20	0.8%	
Portugal	19	0.8%	

COUNTRY	POSTINGS	SHARE	REGION
Ireland	19	0.8%	

Table 1 — Top 15 countries by unique 2026 AI Product Manager postings.

The city picture

City-level concentration is sharper than country-level concentration. San Francisco alone carries 6.7% of global postings, and the top five cities — San Francisco, San Jose, New York, Bengaluru and London — together account for 24.2%. These counts exclude roles posted as remote-first with no anchor city, so they understate distributed hiring.

The cities hiring the most AI Product Managers

Unique 2026 postings by primary posting city (remote-first roles excluded from city counts)

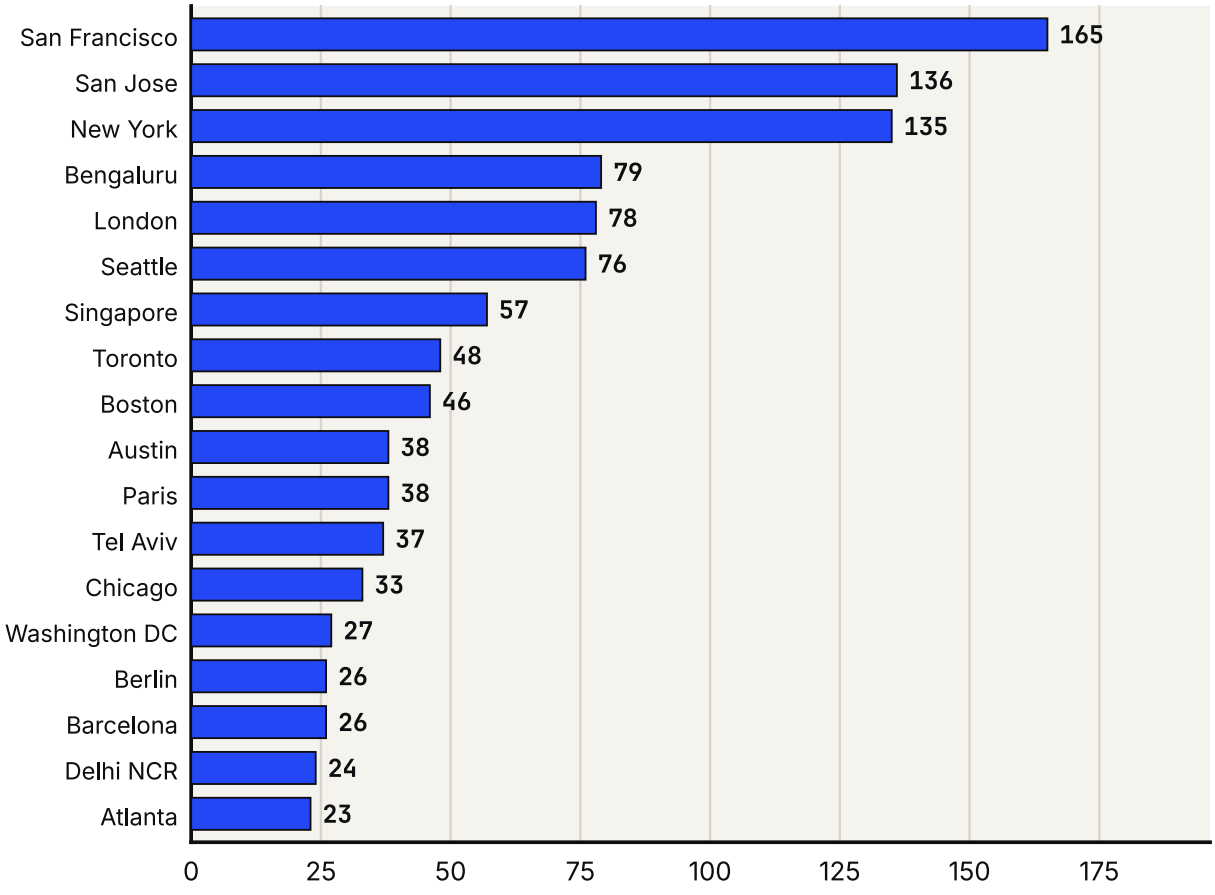

San Francisco leads with 165 postings — 6.7% of the global sample.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 3 — The cities posting the most AI Product Manager roles.

Which industries are buying

Enterprise Software & SaaS is the largest single sector at 18%, followed by Financial Services & FinTech (17%) and Healthcare & Life Sciences (8%). The more useful observation is how wide the spread is: 11 distinct sectors each account for at least 2% of postings. AI product management has already left the technology industry.

Which industries hire AI Product Managers

Share of unique 2026 postings by employer industry, inferred from posting content

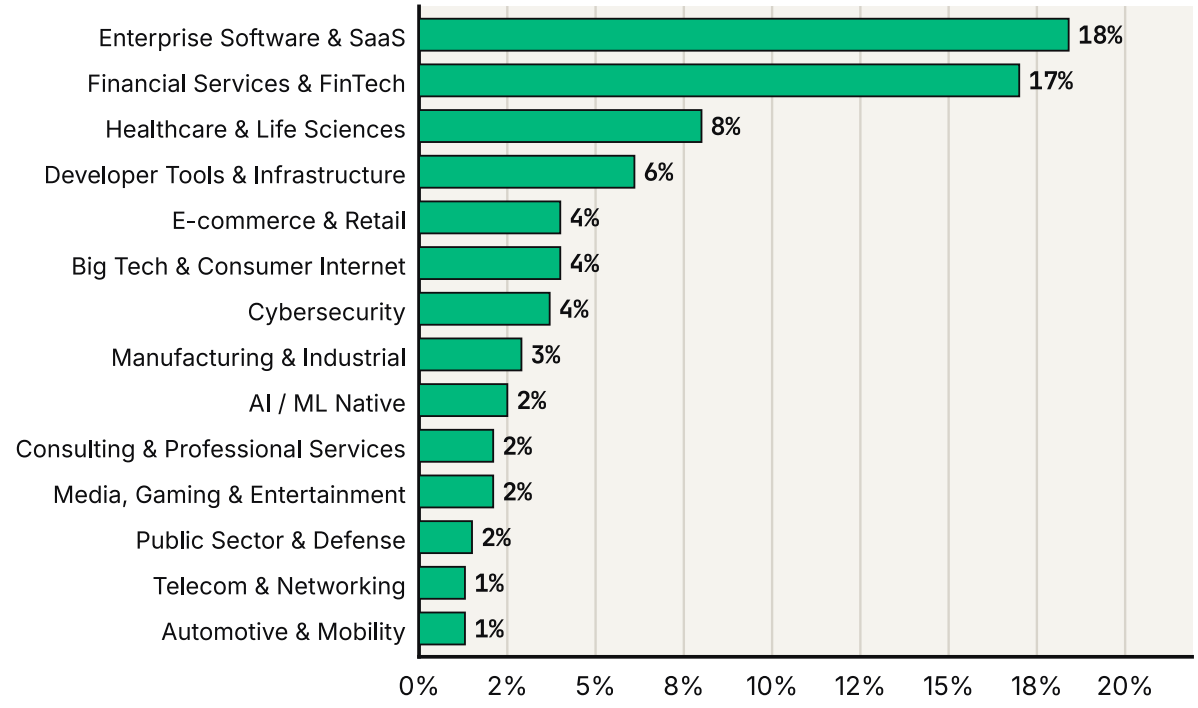

Enterprise Software & SaaS accounts for 18% of AI PM hiring, 1.1x the next sector (Financial Services & FinTech, 17%).

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Industry assigned by weighted keyword classification of the job description.

Figure 4 — Share of AI Product Manager postings by employer industry.

WHAT THE DATA SAYS

Regulated and operationally complex industries — healthcare, financial services, manufacturing, public sector — collectively account for a substantial share of demand. These employers consistently ask for different things than software companies do: governance, evaluation, auditability and domain fluency appear far more often than named frameworks.

Who is posting

The employers posting the most AI Product Manager roles

Unique 2026 postings per company after de-duplication of syndicated and re-posted listings

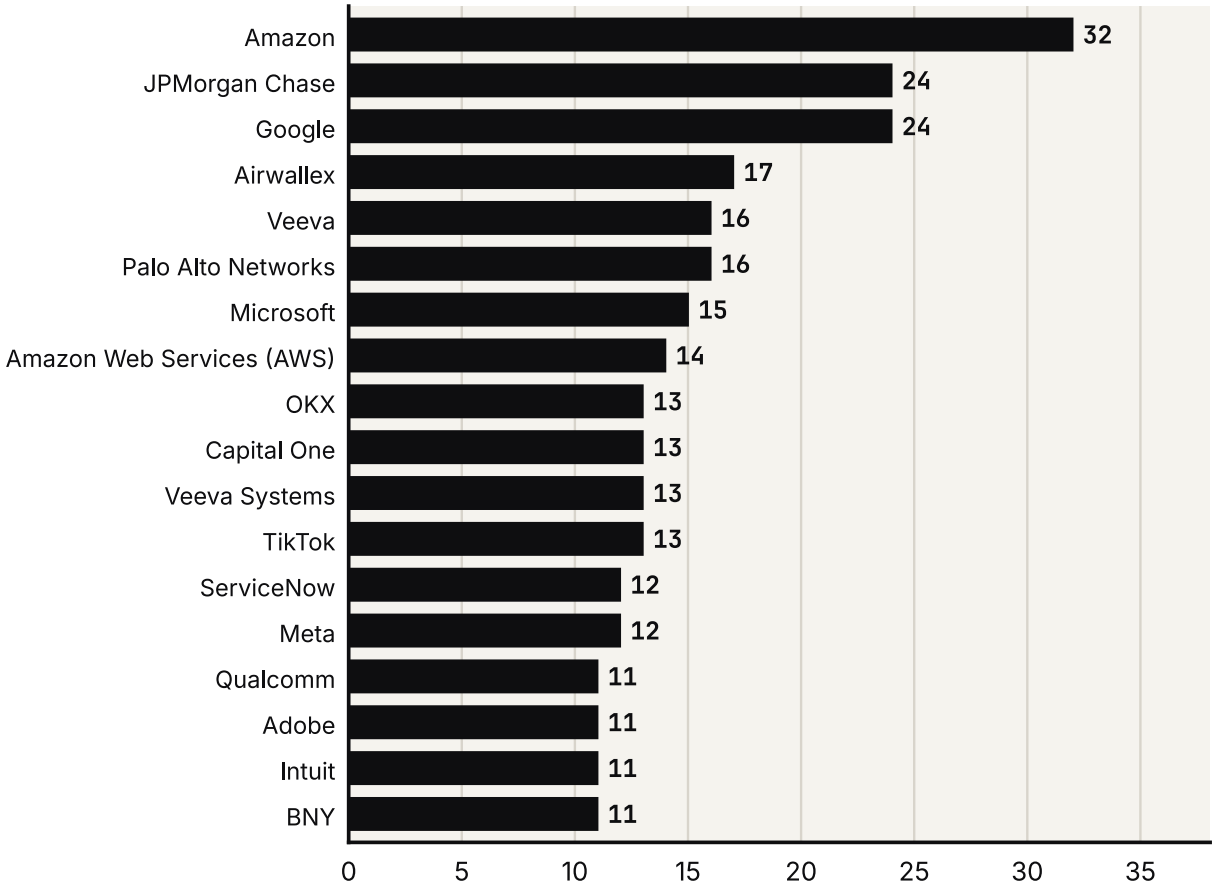

Hiring is broad, not concentrated: 76% of employers posted exactly one AI PM role, and the top 10 account for just 8% of the sample.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Top 10 employers account for 7.5% of the sample; 76% of hiring companies posted exactly one AI PM role.

Figure 5 — Employers with the most unique AI Product Manager postings.

EMPLOYER	POSTINGS	PRIMARY INDUSTRY	COUNTRIES
Amazon	32	E-commerce & Retail	4
JPMorgan Chase	24	Financial Services & FinTech	3

EMPLOYER	POSTINGS	PRIMARY INDUSTRY	COUNTRIES
Google	24	Big Tech & Consumer Internet	3
Airwallex	17	Financial Services & FinTech	4
Veeva	16	Healthcare & Life Sciences	4
Palo Alto Networks	16	Cybersecurity	2
Microsoft	15	Big Tech & Consumer Internet	2
Amazon Web Services (AWS)	14	Enterprise Software & SaaS	1
OKX	13	Financial Services & FinTech	3
Capital One	13	Financial Services & FinTech	1
Veeva Systems	13	Healthcare & Life Sciences	3
TikTok	13	Big Tech & Consumer Internet	2
ServiceNow	12	Enterprise Software & SaaS	2
Meta	12	Big Tech & Consumer Internet	3

Table 2 — Employers ranked by unique postings after de-duplication of syndicated listings.

A note on what this dataset can say about momentum

Monthly counts across 2026 year-to-date run from 12 to 786 postings by month of first publication. It is tempting to read that as explosive growth. It is not, and this report will not claim it is.

A job advertised in January and filled in March is no longer visible in September. Every snapshot-based dataset therefore shows a rising line toward the capture date, regardless of what hiring actually did. The chart below is included because it describes how fresh this sample is — a methodological fact readers should know — not because it measures demand over time.

How recent the dataset is, by month of first publication

Unique 2026 postings still live when the data was captured, by the month each was first published

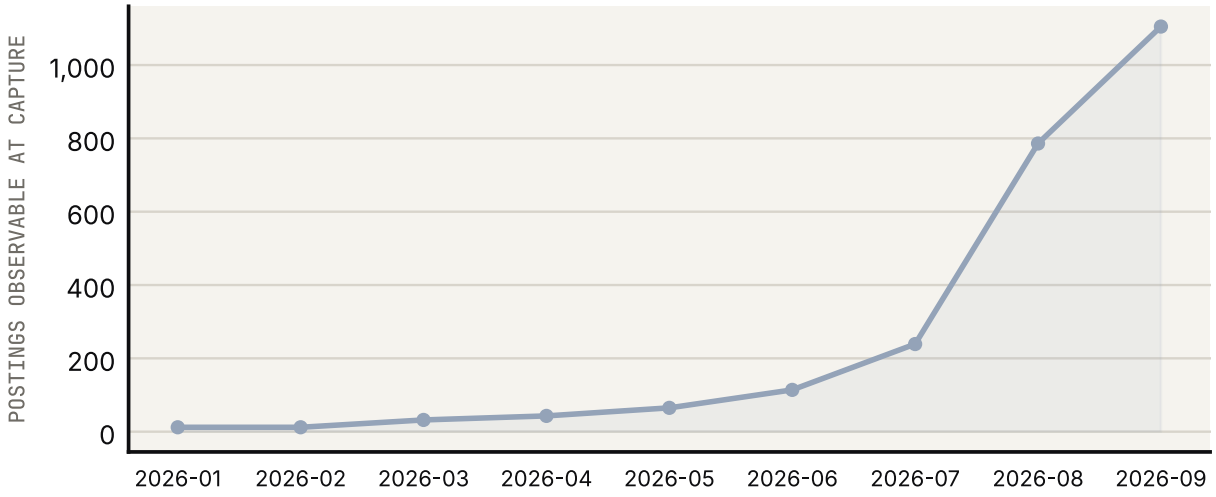

The upward slope is survivorship, not growth: older postings have been filled and de-listed. This chart says how fresh the sample is, and nothing about hiring trend.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. This is a recency profile of the sample, not a measure of hiring volume. A job posted in January and filled in March is invisible at a September capture, so early months are systematically under-represented. The final month is also partial.

Figure 6 — Recency profile of the dataset: unique 2026 postings still observable at capture, by month of first publication.

READ THIS CAREFULLY

Do not cite this chart as evidence of growth in AI Product Manager hiring. It measures survivorship, not demand. A defensible trend measurement requires repeated captures using one method; the 2026 baseline established in this report is the intended comparison point for a 2027 edition.

SECTION 03

03 What an AI Product Manager actually is

Titles, archetypes, responsibilities, and the real differences from traditional product management.

1,808 distinct job titles in the sample

73% carry AI, ML or GenAI explicitly in the title

Seven recurring role archetypes with distinct skill profiles

There is no single AI Product Manager job. There are seven of them, and the differences between them matter more to a candidate's preparation than anything in the title.

The sample contains 1,808 distinct raw job titles. 73% name AI, ML, GenAI or a related term in the title itself (Tier A in this report's classification). The remainder are roles titled as ordinary product management jobs where the description makes AI a central product responsibility (Tier B). Both are counted; the distinction is preserved throughout so readers can see how much of a finding depends on it.

The most common titles

JOB TITLE (AS POSTED)	POSTINGS	SHARE
AI Product Manager	103	4.2%
Senior Product Manager	50	2.0%
Product Manager	39	1.6%
AI Product Owner	37	1.5%
Senior AI Product Manager	31	1.3%
Senior Product Manager, AI	17	0.7%
AI Product Lead	16	0.7%
Principal Product Manager	13	0.5%
Product Manager, AI	12	0.5%
Senior Product Manager - AI	11	0.4%
Staff Product Manager	9	0.4%

JOB TITLE (AS POSTED)	POSTINGS	SHARE
Technical Product Manager	9	0.4%
Product Manager - AI	9	0.4%
Principal Product Manager, AI	6	0.2%
Lead AI Product Manager	6	0.2%
Staff Product Manager, Developer Experience	5	0.2%

Table 3 — Most frequent raw job titles in the sample.

WHAT THE DATA SAYS

Title inflation is real but shallow. The most common titles are the plainest ones. Exotic titles — prompt engineer-adjacent, 'AI evangelist', 'agent architect' — are statistically negligible in product-management hiring.

Seniority mix

The market is weighted to the middle and upper-middle of the ladder. Mid-level roles are the largest single block at 41%. Entry and associate-level roles account for 2% — the clearest quantitative statement of how hard this field is to enter directly.

Seniority mix of AI Product Manager roles

Share of unique 2026 postings by seniority level parsed from the job title

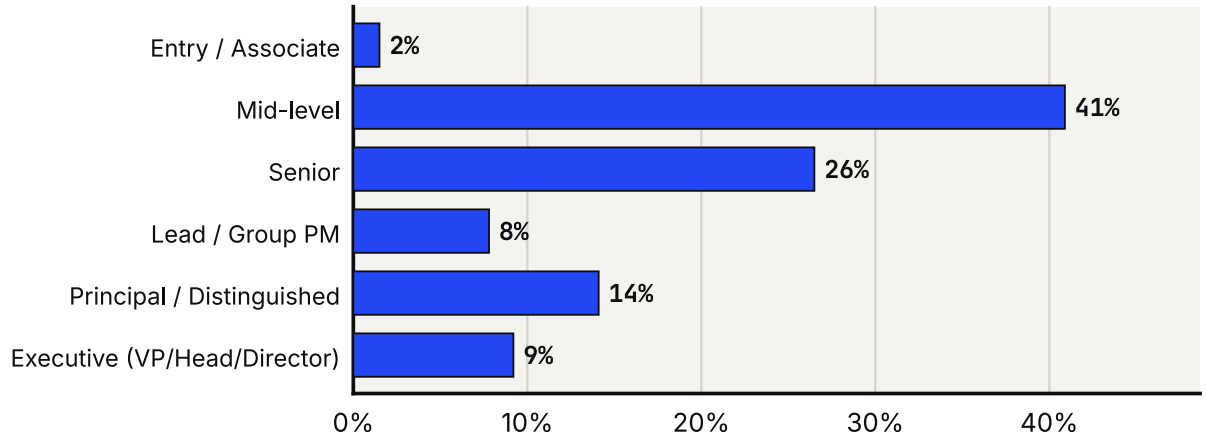

Mid-level roles dominate at 41% — this is a mid-to-senior market.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 7 — Seniority distribution of AI Product Manager postings.

The seven archetypes

Clustering titles and responsibility language produces seven recurring role types. They differ in what they build, who they work with, and what they screen for.

The seven faces of the AI Product Manager

Share of unique 2026 postings by role archetype, assigned from title and responsibility language

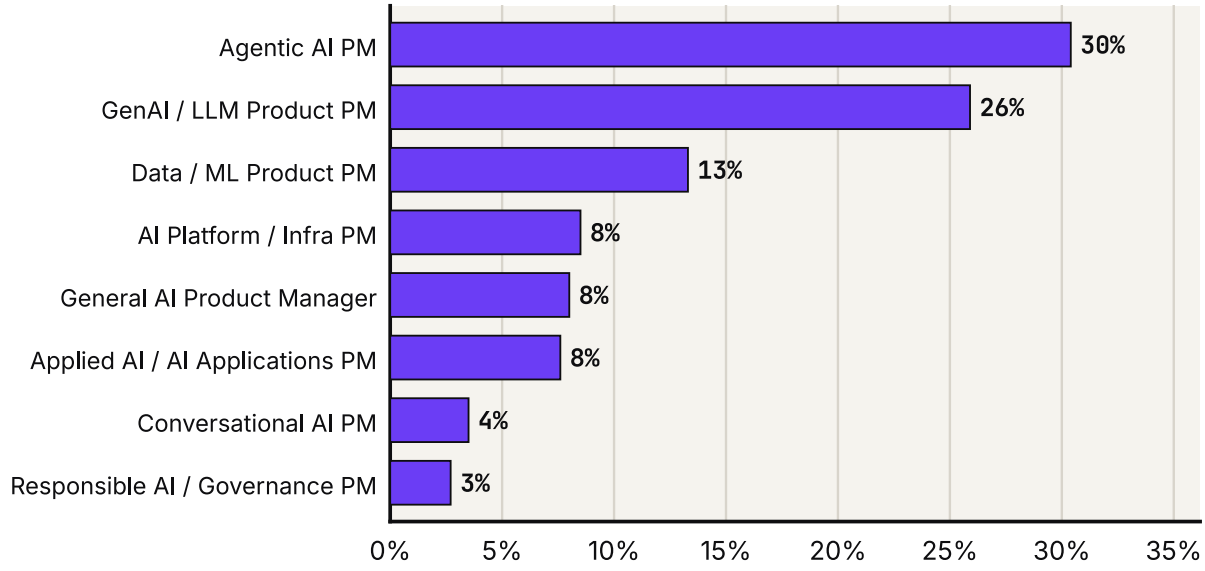

Agentic AI PM is the largest archetype at 30% of the market.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 8 — Share of postings by AI Product Manager role archetype.

Agentic AI PM

30% of postings

Owns autonomous or semi-autonomous workflows. Screens for tool-use design, failure recovery, permissioning and the organisational question of what a system is allowed to do unsupervised.

Most-requested: Artificial Intelligence, AI Agents / Agentic AI, Roadmapping, Communication, Prioritization · median experience 5 yrs · median posted pay $195K

GenAI / LLM Product PM

26% of postings

Owns user-facing products built on foundation models. Screens for prompt and context design, evaluation discipline, latency and cost trade-offs, and comfort with non-deterministic output.

Most-requested: Artificial Intelligence, Large Language Models, Roadmapping, Communication, Collaboration · median experience 5 yrs · median posted pay $185K

Data / ML Product PM

13% of postings

The oldest archetype. Owns ranking, recommendation, forecasting or personalisation. Screens for metrics rigour, experimentation and data-pipeline literacy.

Most-requested: Artificial Intelligence, Roadmapping, Communication, Prioritization, Collaboration · median experience 5 yrs · median posted pay $195K

AI Platform / Infra PM

8% of postings

Owns the internal platform other teams build on: model serving, inference cost, tooling, guardrails. Screens for systems thinking and developer empathy more than consumer product instinct.

Most-requested: Artificial Intelligence, Roadmapping, Prioritization, Communication, Collaboration · median experience 5 yrs · median posted pay $204K

General AI Product Manager

8% of postings

Broad AI ownership without a named specialism, most common at companies adding their first AI PM. Screens for range.

Most-requested: Artificial Intelligence, Roadmapping, Collaboration, Communication, Stakeholder Management · median experience 5 yrs · median posted pay $194K

Applied AI / AI Applications PM

8% of postings

Embeds AI into an existing product surface. Screens for prioritisation and change management as much as AI depth.

Most-requested: Artificial Intelligence, Roadmapping, Communication, Prioritization, Collaboration · median experience 5 yrs · median posted pay $175K

Conversational AI PM

4% of postings

Owns assistants, copilots and support automation. Screens for dialogue design, containment and deflection metrics, and integration with human escalation.

Most-requested: Artificial Intelligence, Roadmapping, Stakeholder Management, Communication, Prioritization · median experience 5 yrs · median posted pay $187K

Responsible AI / Governance PM

3% of postings

Owns policy, safety, evaluation standards and regulatory readiness. Screens for governance frameworks, risk language and cross-functional authority.

Most-requested: Artificial Intelligence, Roadmapping, Responsible AI / AI Ethics, Stakeholder Management, Prioritization · median experience 5 yrs · median posted pay $178K

How AI product management differs from traditional product management

The data supports four concrete differences, and contradicts several assumed ones.

1. **The unit of uncertainty changes.** A traditional PM ships a deterministic feature and measures adoption. An AI PM ships a probabilistic capability and has to decide what "good enough" means before launch. This is why evaluation language appears in postings where no other technical term does.
2. **Quality becomes a product surface, not a QA step.** Guardrails, hallucination handling, human-in-the-loop review and escalation paths appear as product requirements rather than engineering details.
3. **Cost and latency become roadmap items.** Inference cost per request sits in AI PM postings the way infrastructure cost never sat in consumer PM postings.
4. **The partner set widens.** AI PM postings name applied scientists, ML engineers, data engineers, legal and risk functions far more often than conventional product postings do.

> **WHAT THIS MEANS**
>
> What does **not** change is the core of the job. Product Strategy (53%), Roadmapping (76%), Stakeholder Management (64%) and Prioritization (67%) are requested at rates that exceed almost every AI-specific skill in the dataset. Candidates who treat AI PM as a different profession rather than an extension of product management consistently prepare for the wrong interview.

8%

of postings that state a figure accept two years of experience or fewer

This is the single hardest structural fact in the dataset for anyone trying to enter AI product management directly. The field is staffed almost entirely by people moving across from adjacent roles, not by people hired into it from outside.

PART II

The Requirements

Experience, education and the full skill stack, measured from what employers wrote down rather than from what the field says about itself.

04 Experience

05 Degrees

06 The skill stack

07 Most requested skills

08 The modern AI stack

09 Technical depth

SECTION 04

04 The experience question

How much experience employers ask for, how that varies by level and geography, and what it means for people trying to enter the field.

Median requirement: 5 years

8% of postings accept 2 years or fewer

67% of postings state an explicit figure

The median AI Product Manager posting asks for 5 years of experience. That single number explains more about who gets hired into this field than any skill list.

Of 2,448 postings, 1,640 (67%) state an explicit minimum. Across those, the median is 5 years, the mean 5.9 years, the 25th percentile 4 years and the 75th percentile 8 years. The distribution is tight: this is a market that has settled on a fairly narrow band.

How many years of experience employers ask for

Distribution of minimum required years of experience (n=1,640 postings stating a figure)

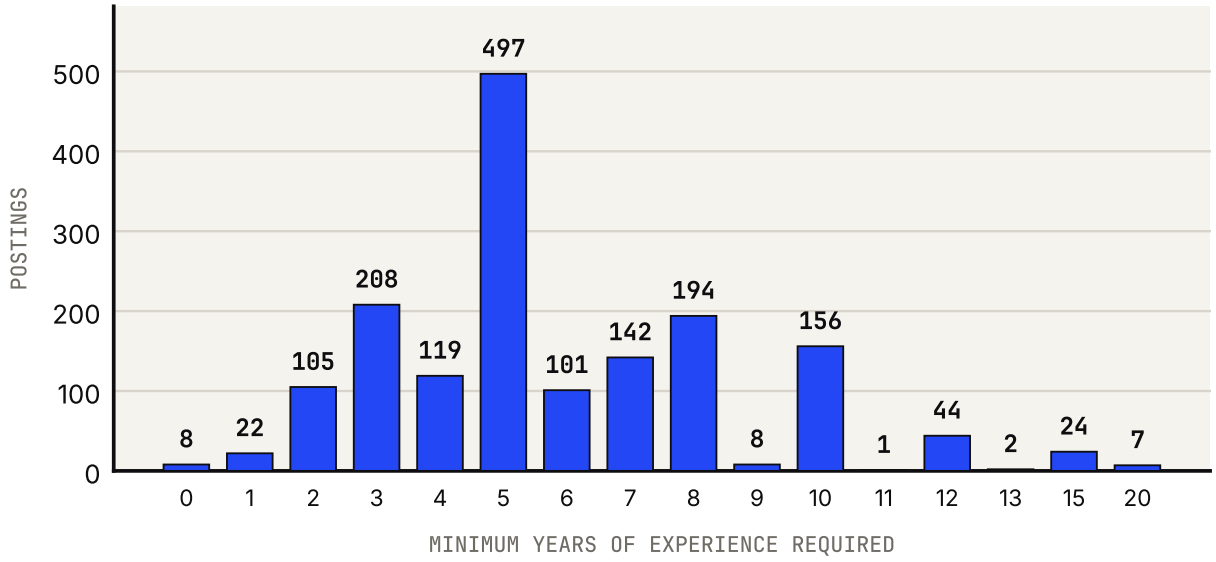

The median AI PM posting asks for 5 years; only 8% accept 2 years or fewer.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. 67% of postings state a minimum experience requirement.

Figure 9 — Distribution of minimum required experience across postings stating a figure.

EXPERIENCE BAND	POSTINGS	SHARE OF POSTINGS STATING A FIGURE
0-2 years	135	8.2%
3-5 years	824	50.2%
6-7 years	243	14.8%
8-9 years	202	12.3%
10+ years	236	14.4%

Table 4 — Minimum experience requirement bands.

The entry-level reality

8% of postings that state a figure accept two years of experience or fewer. Entry and associate-titled roles are 2% of the whole sample. Both numbers point at the same structural fact: there is no meaningful graduate pipeline into AI product management. The field is staffed almost entirely by people moving across from adjacent roles.

FOR JOB SEEKERS

If you have fewer than three years of professional experience, the realistic path is not an AI PM title. It is an adjacent role at a company that ships AI — associate PM, product analyst, solutions or forward-deployed engineer, technical program manager — followed by an internal move. The data shows the door is internal far more often than external.

Experience by seniority

SENIORITY LEVEL	MEDIAN YEARS REQUIRED	POSTINGS STATING A FIGURE
Entry / Associate	2	17
Mid-level	5	618
Senior	5	455
Lead / Group PM	6	124
Principal / Distinguished	8	264
Executive (VP/Head/Director)	8	162

Table 5 — Median stated experience requirement by seniority band.

The steps between levels are smaller than most candidates expect. The gap between a senior and a principal posting is rarely more than a few years of stated experience — the real difference sits in scope language, which Section 15 examines directly.

Experience by market

COUNTRY	MEDIAN YEARS REQUIRED	POSTINGS STATING A FIGURE
Canada	5	82
Germany	5	18
India	5	136
Israel	5	37
Mainland China	5	16
Netherlands	5	13
Poland	5	12
Singapore	5	31
Spain	5	25
United Kingdom	5	63
United States	5	999
France	4	16

Table 6 — Markets with at least 12 postings stating an experience figure.

SECTION 05

05 Do you need a degree?

What employers actually write about education — required, preferred, or never mentioned — and which fields of study appear.

22% state a degree as required

53% never mention a degree at all

MBA appears in 7.7% of postings

Formal education is the most over-weighted anxiety in AI product management. The postings are unambiguous: skills and shipped work are screened first.

Do employers actually require a degree?

How the degree requirement is expressed across unique 2026 AI PM postings

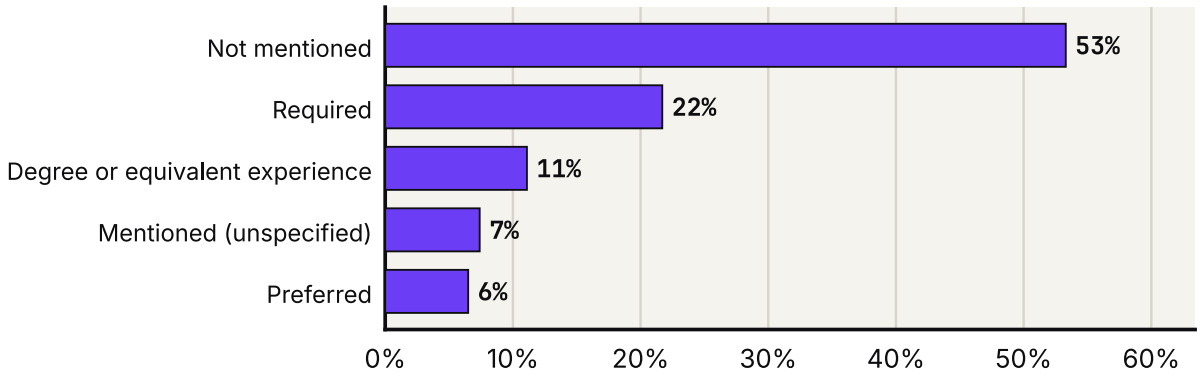

Only 22% state a degree as a hard requirement; 53% never mention one.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. 'Degree or equivalent experience' counts postings that name a degree but explicitly allow an experience substitute.

Figure 10 — How the degree requirement is expressed across AI Product Manager postings.

Only 22% of postings state a degree as a firm requirement. 53% do not mention education at all. A further 11% name a degree but explicitly allow equivalent practical experience in its place — language that has become close to standard in this market.

Which levels and fields appear

DEGREE LEVEL MENTIONED	SHARE OF ALL POSTINGS
Bachelor's	35.7%
Master's	22.8%
MBA	7.7%
PhD	2.3%

Table 7 — A posting may mention more than one level; shares do not sum to 100%.

Where a field of study is named, Computer Science / Engineering leads at 26.5% of all postings, followed by Business / Management (18.8%).

FIELD OF STUDY NAMED	SHARE OF ALL POSTINGS
Computer Science / Engineering	26.5%
Business / Management	18.8%
Data Science / Statistics / Math	10.0%
Engineering (general)	2.8%
Design / HCI	1.6%
Physical / Natural Sciences	1.1%
AI / Machine Learning	0.2%

Table 8 — Fields of study named anywhere in the education requirement.

Degrees by seniority

Degree requirements by seniority

Share of postings at each level, 2026 unique postings

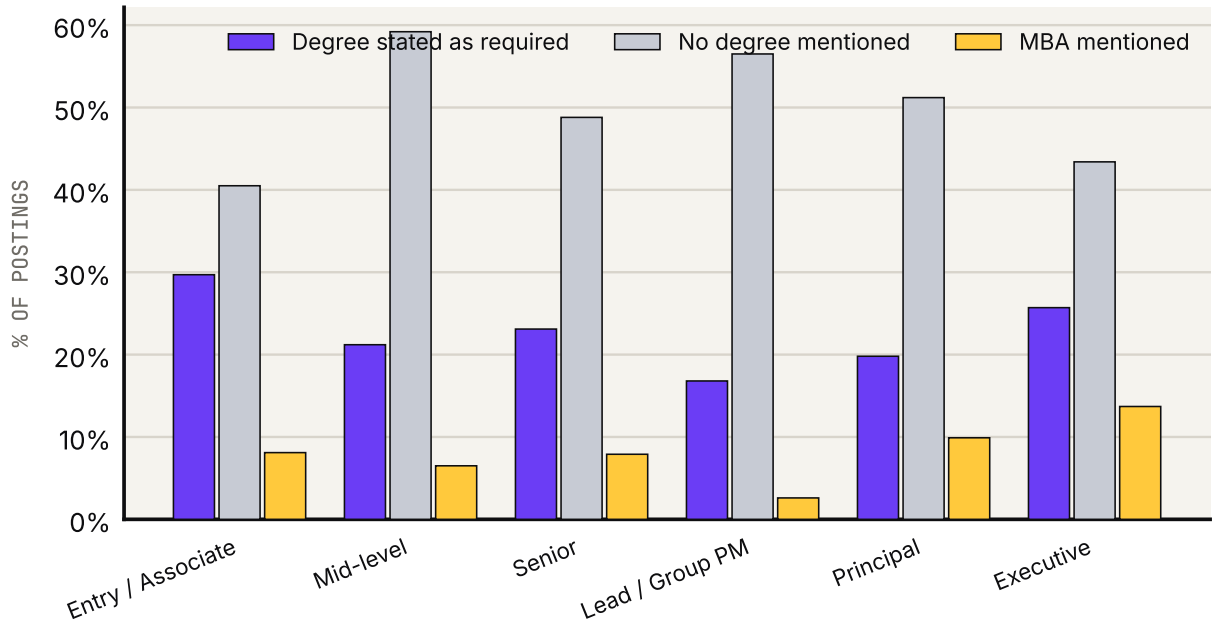

Degree-as-requirement runs from 17% to 30% across levels; MBA mentions peak at 14% (Executive).

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 11 — Degree requirements and MBA mentions by seniority level.

SENIORITY	DEGREE REQUIRED	NO DEGREE MENTIONED	MBA MENTIONED
Entry / Associate	29.7%	40.5%	8.1%
Mid-level	21.2%	59.2%	6.5%
Senior	23.1%	48.8%	7.9%
Lead / Group PM	16.8%	56.5%	2.6%
Principal / Distinguished	19.8%	51.2%	9.9%
Executive (VP/Head/Director)	25.7%	43.4%	13.7%

Table 9 — Education signals by seniority band.

Does an MBA help?

An MBA is named in 7.7% of postings and a PhD in 2.3%. Neither is a gate. Where an MBA does appear it clusters in specific contexts rather than across the market — a signal about the type of employer, not about the profession.

Degree language and posted pay

DEGREE LANGUAGE IN POSTING	POSTINGS DISCLOSING PAY	MEDIAN POSTED SALARY (USD)
Degree or equivalent experience	127	$194,612
Mentioned (unspecified)	61	$195,000
Not mentioned	422	$195,000
Preferred	72	$197,750
Required	222	$185,756

Table 10 — Association only. Degree language correlates with employer type and seniority, both of which move pay independently.

Certifications

Certifications are mentioned in 12.3% of postings. Where named, the most common are Scrum (CSPO/CSM/PSPO) (2.5%), AI/ML certification (0.6%), PMP / PMI (0.6%) and Security (CISSP/CISM) (0.4%).

CERTIFICATION FAMILY	SHARE OF ALL POSTINGS
Scrum (CSPO/CSM/PSPO)	2.50%
AI/ML certification	0.60%
PMP / PMI	0.60%
Security (CISSP/CISM)	0.40%
SAFe	0.40%
Pragmatic Institute	0.40%
Six Sigma / Lean	0.20%
Cloud certification (AWS/Azure/GCP)	0.00%

Table 11 — Certification families named in postings.

FOR JOB SEEKERS

No certification in this dataset appears often enough to function as a hiring signal. Time spent on a certificate is almost always better spent on a shipped artefact that demonstrates the same skill — Section 21 lists the projects the demand data actually supports.

SECTION 06

06 The AI Product Manager skill stack

The six layers of capability employers assemble when they write an AI Product Manager job description, and how much weight each carries.

Six layers, measured by request frequency

Product layer is requested more often than the AI layer

140-term normalised ontology; synonyms merged

Read enough AI PM job descriptions and the same six layers appear in almost all of them. What varies is the depth demanded at each layer — not which layers exist.

The layers below are constructed from the normalised skill ontology used throughout this report. Percentages are the share of postings requesting at least one skill in that layer.

Foundation layer — how you work

top skill 71%

Communication, collaboration, working through ambiguity, influence without authority. Requested at near-universal rates and the most common reason a strong technical candidate is rejected.

Most requested: Communication 71%, Collaboration 67%, Ambiguity / Autonomy 53%, Leadership 51%, Mentorship / Coaching 21%

Product layer — the actual job

top skill 76%

Strategy, discovery, roadmapping, requirements, prioritisation, experimentation, go-to-market. This is the layer employers screen hardest and the layer candidates most often under-prepare.

Most requested: Roadmapping 76%, Prioritization 67%, Stakeholder Management 64%, Cross-functional Leadership 57%, Product Strategy 53%

AI layer — fluency, not research

top skill 99%

Understanding what models can and cannot do, how they fail, how they are evaluated, and what it costs to run them. Depth of vocabulary matters more than depth of mathematics.

Most requested: Artificial Intelligence 99%, Machine Learning 52%, Large Language Models 46%, Generative AI 29%, Natural Language Processing 5%

Technical layer — enough to be dangerous

top skill 28%

APIs, data, system design, cloud concepts, and enough Python or SQL to interrogate your own product. Requested widely, but rarely at production-engineering depth.

Most requested: REST APIs 28%, APIs 28%, Data Pipelines / ETL 13%, System Design / Architecture 11%, SQL 11%

Business layer — why it gets funded

top skill 24%

Business cases, pricing, ROI, competitive analysis, unit economics. The layer that separates a feature owner from a product owner.

Most requested: Partnerships / BD 24%, Sales Enablement 16%, Business Strategy 15%, P&L / Business Metrics 15%, Pricing & Monetization 12%

Leadership layer — scope beyond yourself

top skill 71%

Mentoring, executive communication, cross-organisational alignment. Scales sharply with seniority and is the clearest dividing line in the premium pay segment.

Most requested: Communication 71%, Collaboration 67%, Ambiguity / Autonomy 53%, Leadership 51%, Mentorship / Coaching 21%

WHAT THIS MEANS

The ordering is the finding. Product and human skills occupy most of the top of the table; AI-specific capabilities appear high but as **context** for product work rather than as replacements for it. A candidate who can run discovery, write a defensible PRD for a probabilistic system and explain an evaluation plan is addressing more of this table than one who can name every vector database.

The 25 most-requested skills in AI Product Manager postings

Share of unique 2026 postings mentioning each normalised skill

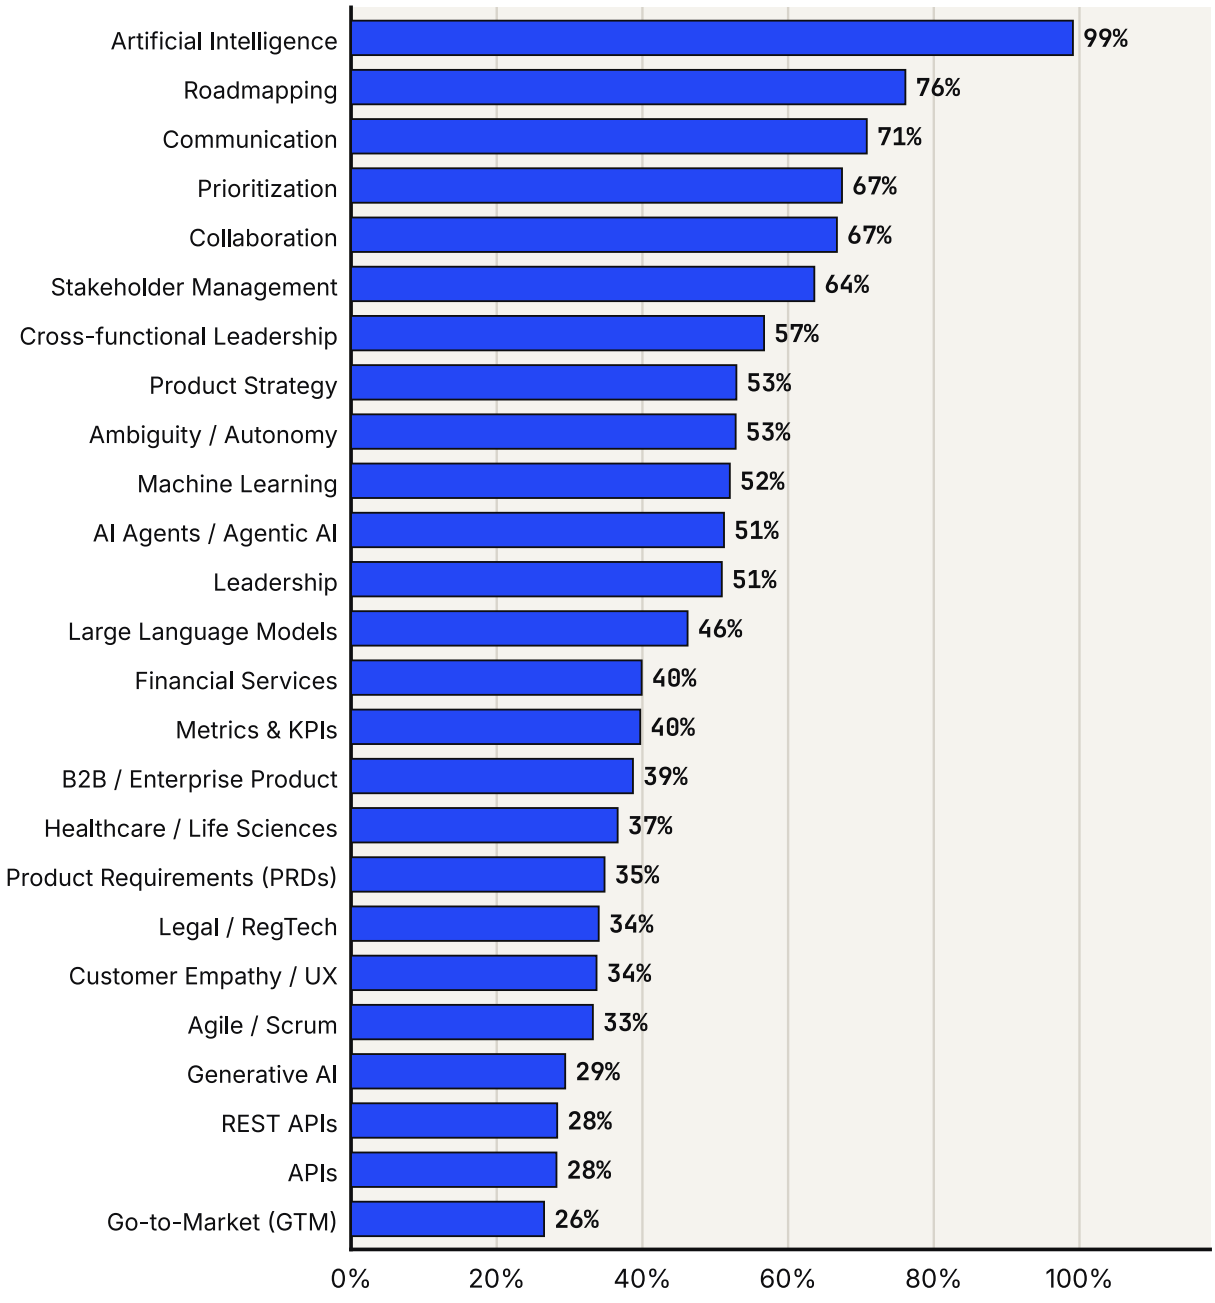

Artificial Intelligence appears in 99% of postings — the single most universal requirement.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Skills normalised across synonyms, abbreviations and plurals using a 140-term ontology.

Figure 12 — The 25 most-requested skills across all AI Product Manager postings.

SECTION 07

07 The most requested skills

Full rankings by category, what appears as a requirement versus a preference, and how the picture changes across markets.

Top skill: Artificial Intelligence at 99%

Required versus preferred measured separately

Cross-market comparison for the largest economies

Frequency alone is a weak guide to what to learn. A skill that appears in 80% of postings is table stakes; a skill that appears in 20% but always in the required block is a filter.

The full top 30

#	SKILL	CATEGORY	% OF POSTINGS	REQUIRED SHARE
1	Artificial Intelligence	AI & ML Foundations	99.1%	68%
2	Roadmapping	Product Management	76.1%	86%
3	Communication	Leadership & Human Skills	70.8%	85%
4	Prioritization	Product Management	67.4%	89%
5	Collaboration	Leadership & Human Skills	66.7%	77%
6	Stakeholder Management	Product Management	63.6%	89%
7	Cross-functional Leadership	Product Management	56.7%	85%
8	Product Strategy	Product Management	52.9%	78%
9	Ambiguity / Autonomy	Leadership & Human Skills	52.8%	80%
10	Machine Learning	AI & ML Foundations	52.0%	73%
11	AI Agents / Agentic AI	Modern GenAI Stack	51.2%	67%
12	Leadership	Leadership & Human Skills	50.9%	81%
13	Large Language Models	AI & ML Foundations	46.2%	74%
14	Financial Services	Domain Knowledge	39.9%	60%
15	Metrics & KPIs	Product Management	39.7%	88%
16	B2B / Enterprise Product	Product Management	38.7%	70%
17	Healthcare / Life Sciences	Domain Knowledge	36.6%	56%
18	Product Requirements (PRDs)	Product Management	34.8%	96%

#	SKILL	CATEGORY	% OF POSTINGS	REQUIRED SHARE
19	Legal / RegTech	Domain Knowledge	34.0%	61%
20	Customer Empathy / UX	Product Management	33.7%	84%
21	Agile / Scrum	Product Management	33.2%	82%
22	Generative AI	AI & ML Foundations	29.4%	77%
23	REST APIs	Technical & Engineering	28.3%	63%
24	APIs	Technical & Engineering	28.2%	80%
25	Go-to-Market (GTM)	Product Management	26.5%	90%
26	Product Lifecycle Mgmt	Product Management	25.6%	86%
27	Data-Driven Decisions	Product Management	24.4%	84%
28	Partnerships / BD	Business & Commercial	23.9%	75%
29	Responsible AI / AI Ethics	Modern GenAI Stack	23.2%	67%
30	Mentorship / Coaching	Leadership & Human Skills	21.2%	71%

Table 12 — 'Required share' is the proportion of a skill's section-placed mentions that appear in a required rather than a preferred block.

Required versus preferred

Many job descriptions separate hard requirements from nice-to-haves. Measuring skills against that split is the single most useful way to read the demand data: it distinguishes gates from bonuses.

Required versus preferred: what is a gate and what is a bonus

Share of unique 2026 postings naming each skill in a required or a preferred block

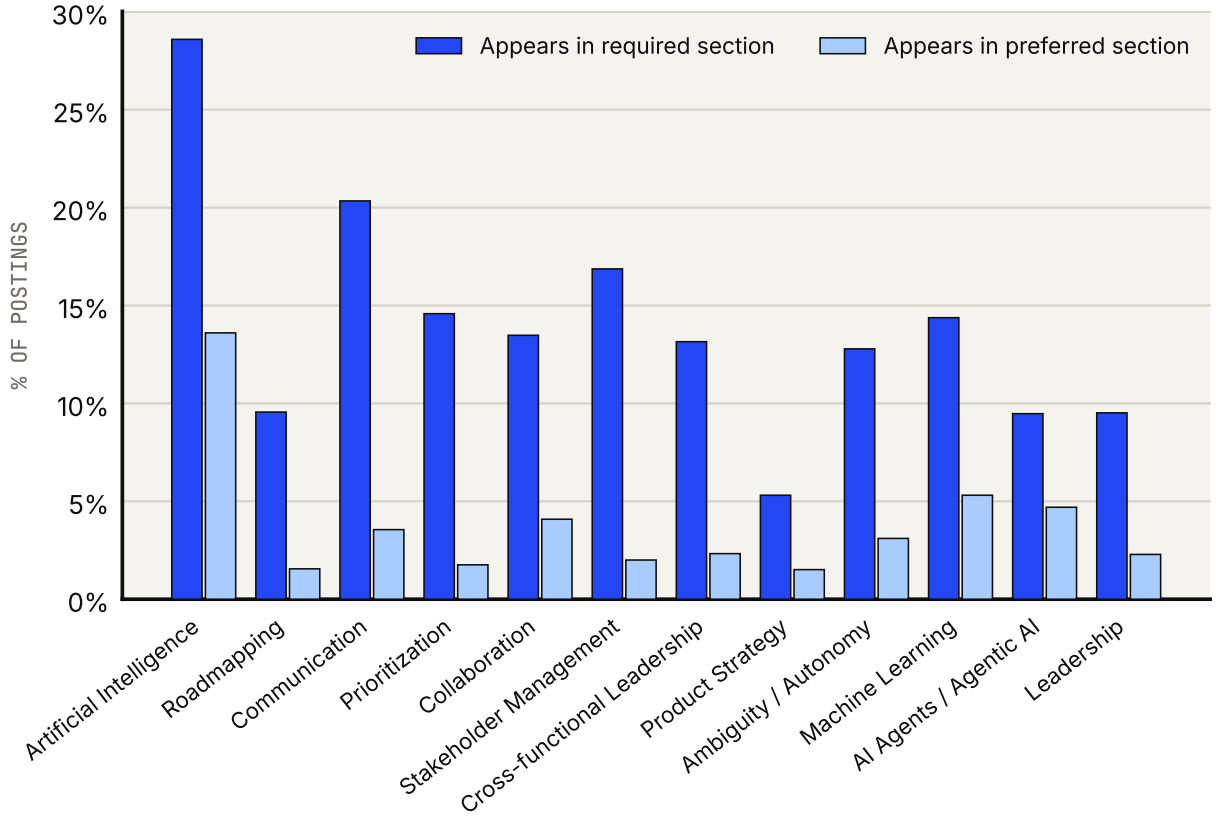

Product Requirements (PRDs) is the most gate-like requirement — 96% of its section-placed mentions sit in a required block — while Healthcare / Life Sciences is the most bonus-like at 56%.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Based on postings whose descriptions separate required and preferred qualifications.

Figure 13 — Where the most common skills appear: required blocks versus preferred blocks.

Skills by category

Product Management

Led by Roadmapping at 76%, ahead of Prioritization (67%) and Stakeholder Management (64%).

Most-requested product-management skills

Share of unique 2026 AI PM postings mentioning each product-management skill

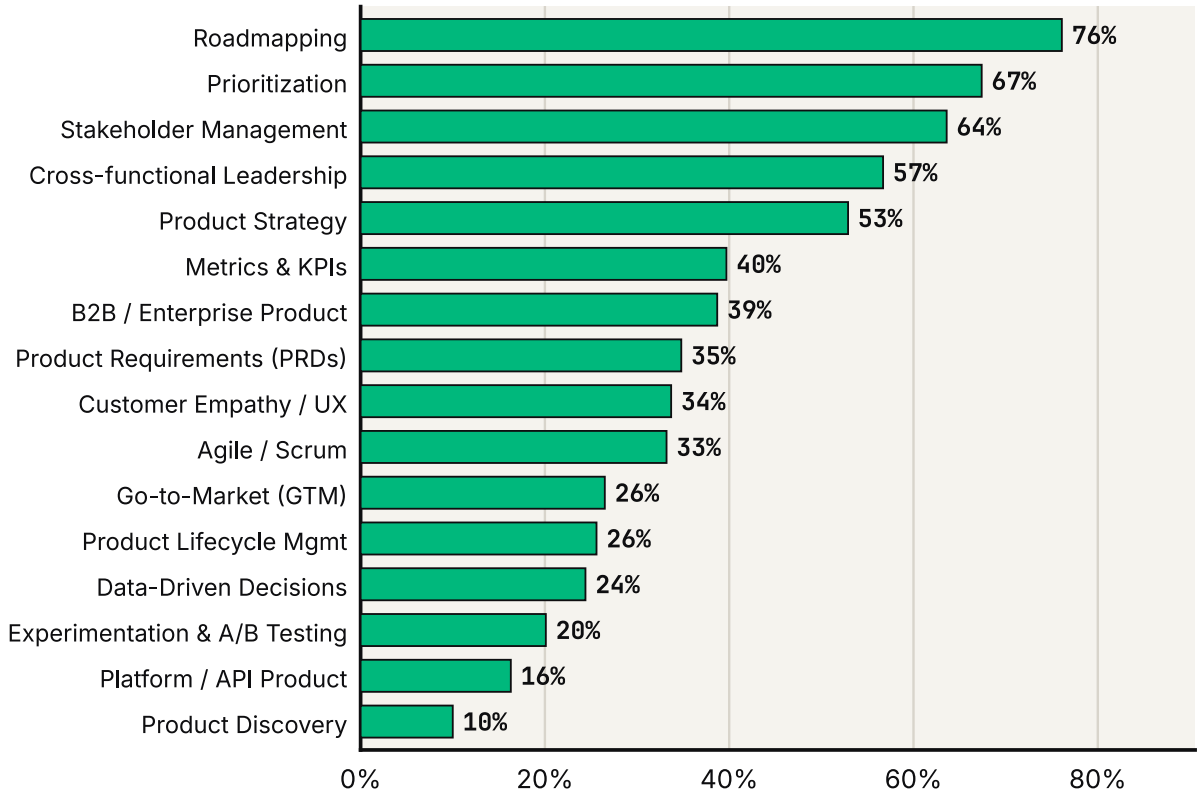

Roadmapping leads at 76% of postings, ahead of Prioritization (67%).

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 14 — Most-requested Product Management skills.

Read this table as the definition of the job. Every item on it predates generative AI and none of them have been displaced by it. The ordering is also instructive: employers ask for roadmapping and prioritisation — the allocation of scarce engineering time — more often than they ask for discovery. That is a statement about what AI product teams are short of. They have more ideas than they can build, and the constraint they are hiring against is judgement about which ones to build. **Also requested in this category:** Data-Driven Decisions 24.4%, Experimentation & A/B Testing 20.1%, Platform / API Product 16.3%, Product Discovery 10.0%, User Research 9.5%, Customer Interviews 6.7%, B2C / Consumer Product 4.7%, Product-Market Fit 3.1%, Technical Documentation 3.0%

AI & ML Foundations

Led by Artificial Intelligence at 99%, ahead of Machine Learning (52%) and Large Language Models (46%).

Most-requested AI and machine-learning skills

Share of unique 2026 AI PM postings mentioning each AI and machine-learning skill

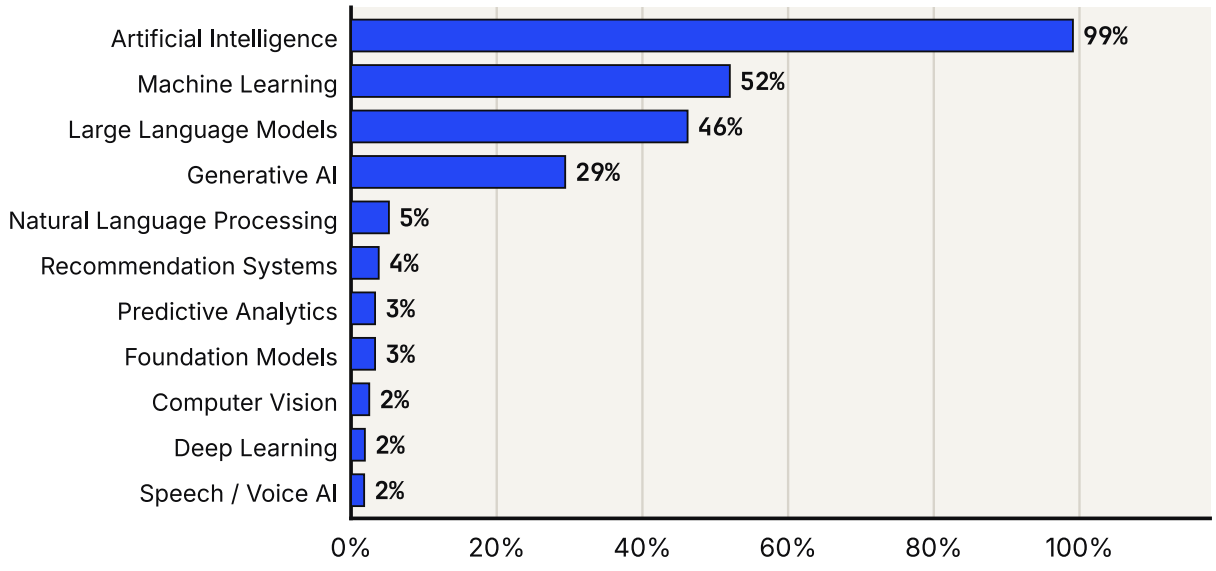

Artificial Intelligence leads at 99% of postings, ahead of Machine Learning (52%).

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 15 — Most-requested AI & ML Foundations skills.

The shape of this table matters more than its top line. Broad terms sit at the top and specific techniques fall away sharply beneath them. Employers are describing a field of work, not a curriculum. A candidate who can explain what a model does, where it breaks and how you would know it is working satisfies most of this table; a candidate who has memorised architectures satisfies very little of it that the first candidate does not already cover.

Technical & Engineering

Led by REST APIs at 28%, ahead of APIs (28%) and Data Pipelines / ETL (13%).

Most-requested technical and engineering skills

Share of unique 2026 AI PM postings mentioning each technical and engineering skill

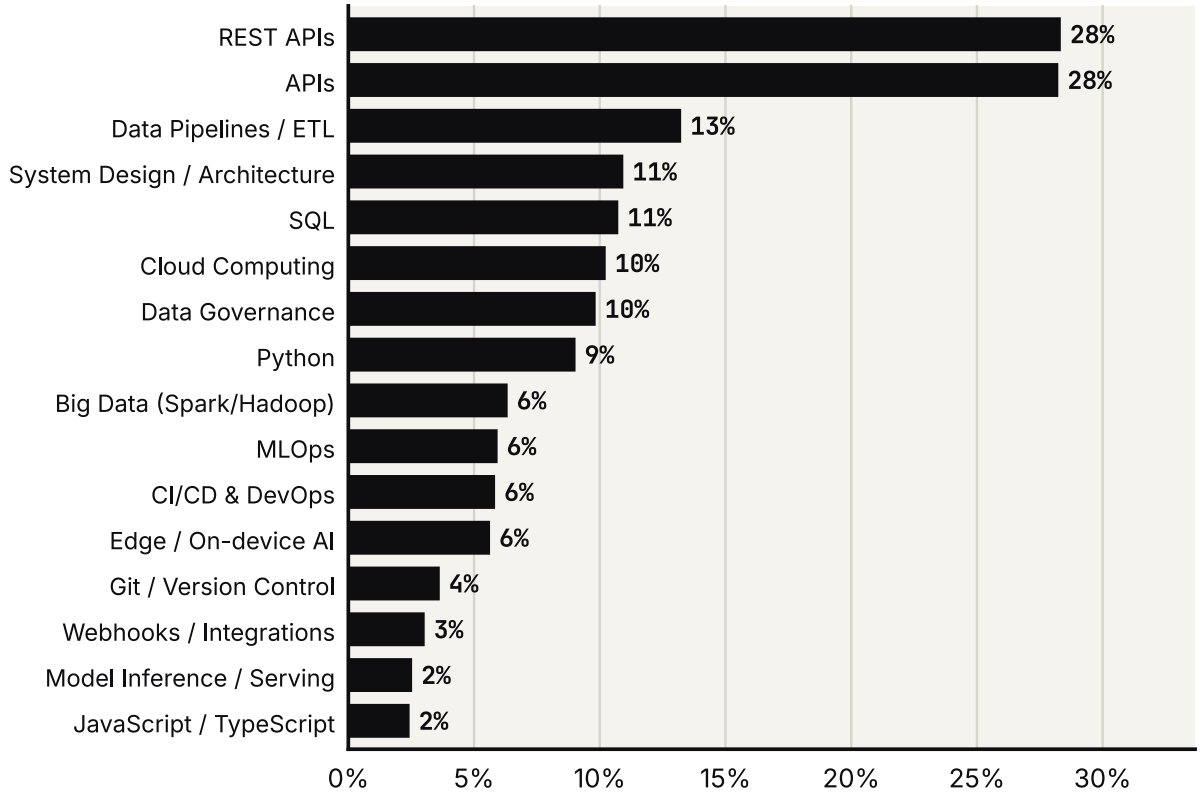

REST APIs leads at 28% of postings, ahead of APIs (28%).

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 16 — Most-requested Technical & Engineering skills.

Note what leads here. APIs and integration language outrank programming languages, and data-pipeline literacy outranks infrastructure tooling. This is the signature of a role that consumes technical systems rather than building them. The practical reading: learn to reason about interfaces, data flow, latency and cost before you learn another language. **Also requested in this category:** Git / Version Control 3.6%, Webhooks / Integrations 3.0%, Model Inference / Serving 2.5%, JavaScript / TypeScript 2.4%, LLMOps 2.2%, Data Modeling 2.1%, Kubernetes 2.1%, Microservices 1.6%, Feature Stores 1.2%, Docker 0.7%

Business & Commercial

Led by Partnerships / BD at 24%, ahead of Sales Enablement (16%) and Business Strategy (15%).

Most-requested business and commercial skills

Share of unique 2026 AI PM postings mentioning each business and commercial skill

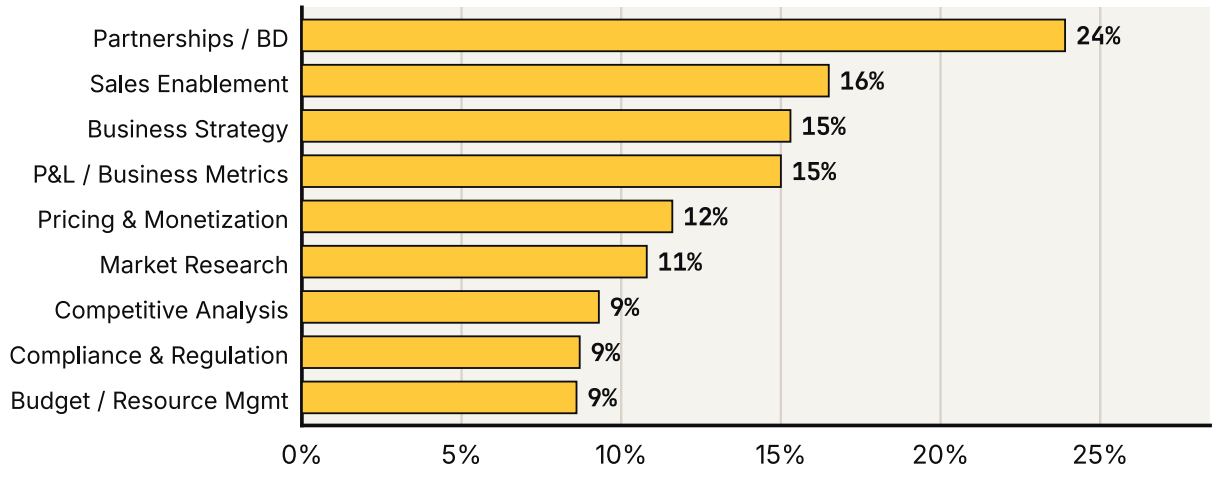

Partnerships / BD leads at 24% of postings, ahead of Sales Enablement (16%).

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 17 — Most-requested Business & Commercial skills.

This is the smallest of the major categories and the most diagnostic. Commercial language appears far more often in senior and higher-paying postings than in mid-level ones, which makes it the clearest leading indicator of scope in the dataset. A candidate who can argue for a budget, price a capability or defend a cost ceiling is describing a different job from one who can only ship features.

Leadership & Human Skills

Led by Communication at 71%, ahead of Collaboration (67%) and Ambiguity / Autonomy (53%).

Most-requested leadership and human skills

Share of unique 2026 AI PM postings mentioning each leadership and human skill

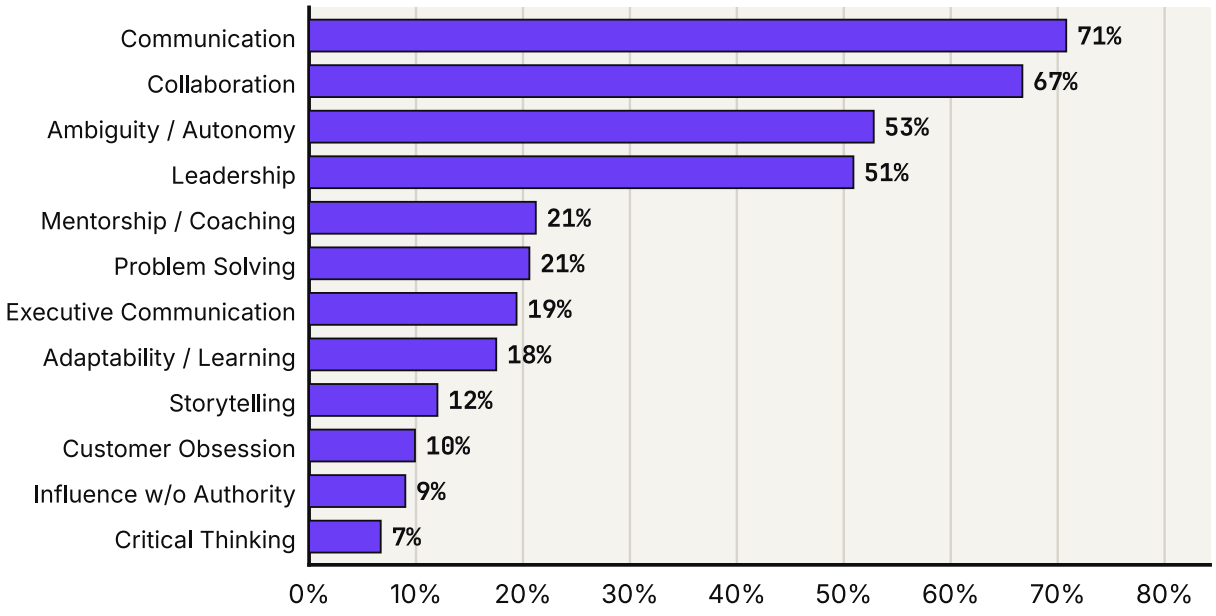

Communication leads at 71% of postings, ahead of Collaboration (67%).

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 18 — Most-requested Leadership & Human Skills skills.

These are the highest-frequency items in the entire ontology, and they are also the ones candidates most often treat as filler. They are not filler. In a role that sits between applied scientists, engineers, legal, risk and commercial teams, the ability to hold four groups to one definition of ‘working’ is the job. Interview processes probe this harder than any technical topic.

Tools & Platforms

Led by Anthropic / Claude at 13%, ahead of Model Context Protocol (MCP) (9%) and Jira / Confluence (8%).

Most-requested tools and platform skills

Share of unique 2026 AI PM postings mentioning each tools and platform skill

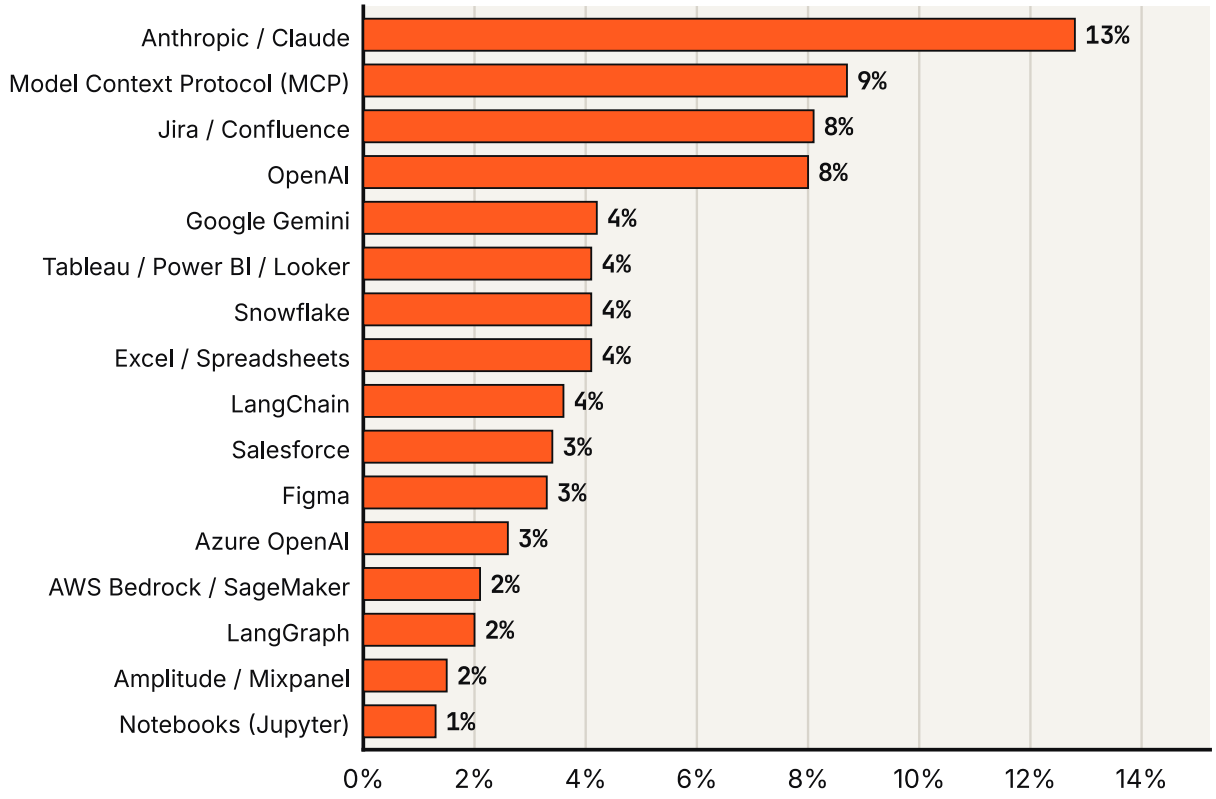

Anthropic / Claude leads at 13% of postings, ahead of Model Context Protocol (MCP) (9%).

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 19 — Most-requested Tools & Platforms skills.

This is the shortest table in the report for a reason. Named vendor tools are the least-requested category in the ontology, and the frequencies here are an order of magnitude below the product and human skills above. Tools churn annually; the concepts underneath them do not. Treat this table as a list of things to be conversant in, never as a study plan. **Also requested in this category:** AWS Bedrock / SageMaker 2.1%, LangGraph 2.0%, Amplitude / Mixpanel 1.5%, Notebooks (Jupyter) 1.3%, PyTorch / TensorFlow 1.2%, LlamaIndex 0.9%, Hugging Face 0.6%, MLflow / W&B 0.4%, Weaviate / Qdrant / Milvus/ FAISS 0.2%, Pinecone 0.2%

Domain Knowledge

Led by Financial Services at 40%, ahead of Healthcare / Life Sciences (37%) and Legal / RegTech (34%).

Most-requested domain-knowledge skills

Share of unique 2026 AI PM postings mentioning each domain-knowledge skill

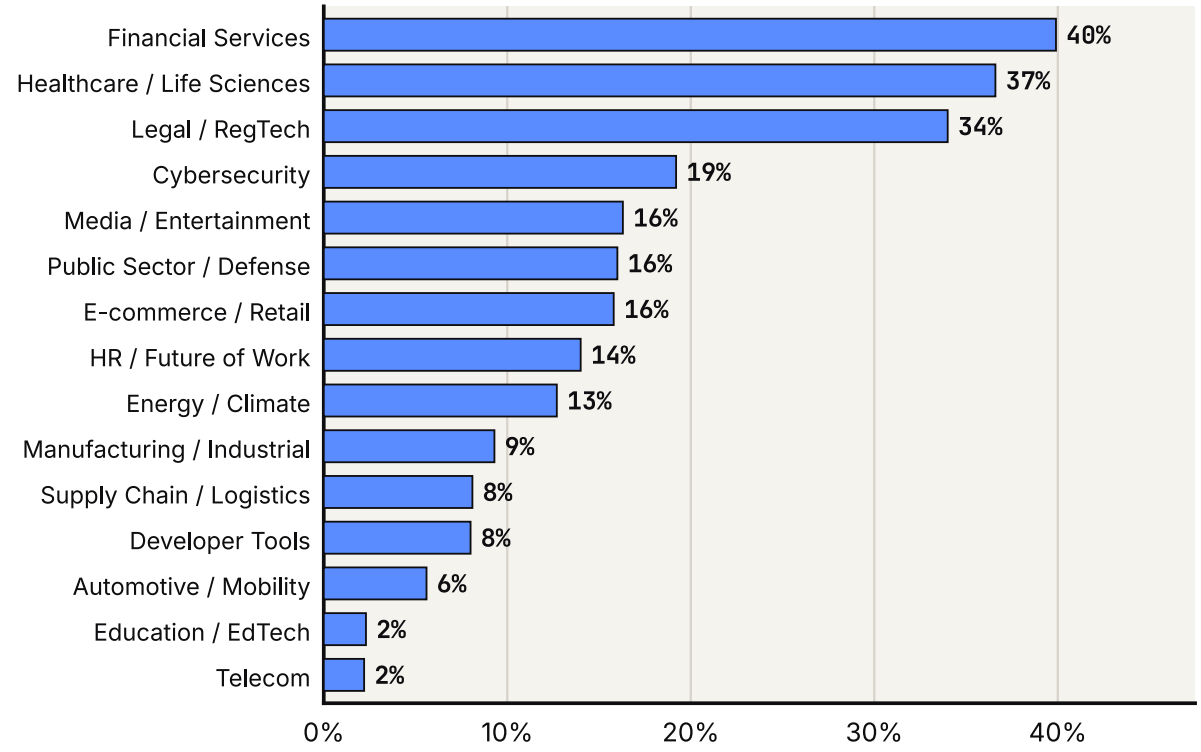

Financial Services leads at 40% of postings, ahead of Healthcare / Life Sciences (37%).

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 20 — Most-requested Domain Knowledge skills.

Domain knowledge behaves differently from every other category: it is rarely a requirement in general postings and frequently decisive in specific ones. Where it appears it is usually non-negotiable, because the employer is buying an understanding of a workflow they cannot teach quickly. For candidates, this is the cheapest differentiator available — you probably already have a domain, and stating it plainly narrows the competition dramatically.

How skills travel together

Skills are not requested independently. The co-occurrence structure below shows which capabilities employers bundle — and therefore which combinations a candidate should build together rather than sequentially.

Which AI PM skills travel together

Of postings requesting the row skill, the percentage that also request the column skill

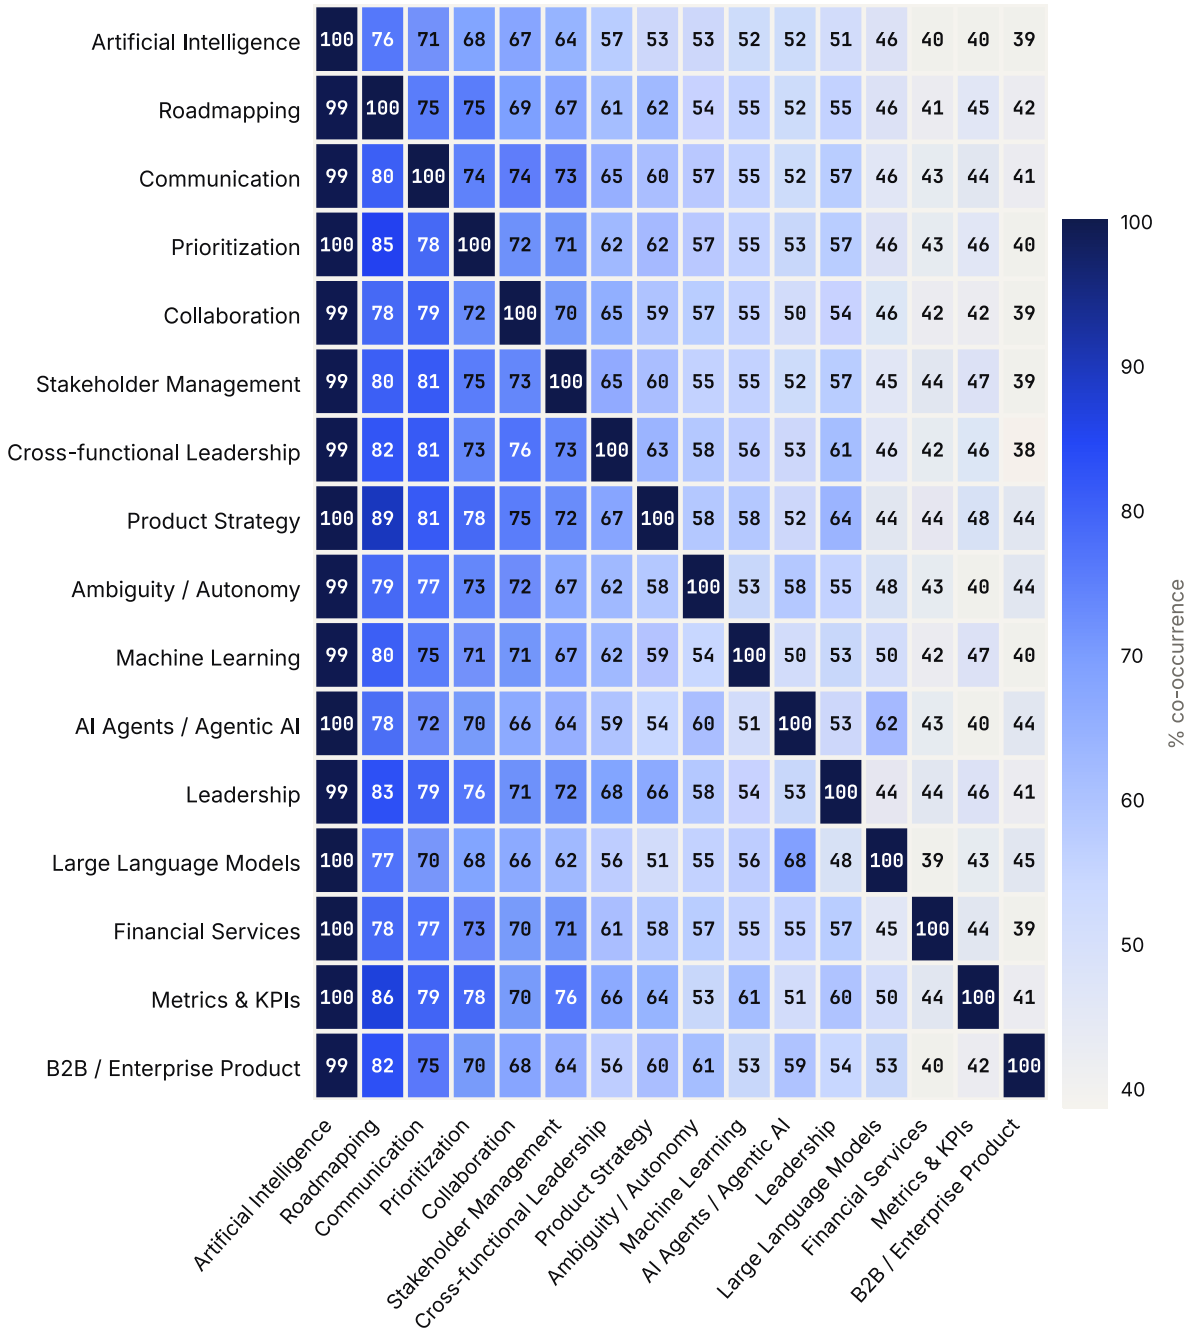

Requirements arrive in bundles: 100% of postings asking for AI Agents / Agentic AI also ask for Artificial Intelligence.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Diagonal is 100% by definition.

Figure 21 — Conditional co-occurrence: of postings requesting the row skill, the share also requesting the column skill.

The AI Product Manager skill network

The 18 most-connected skills; line weight is the number of postings requesting both

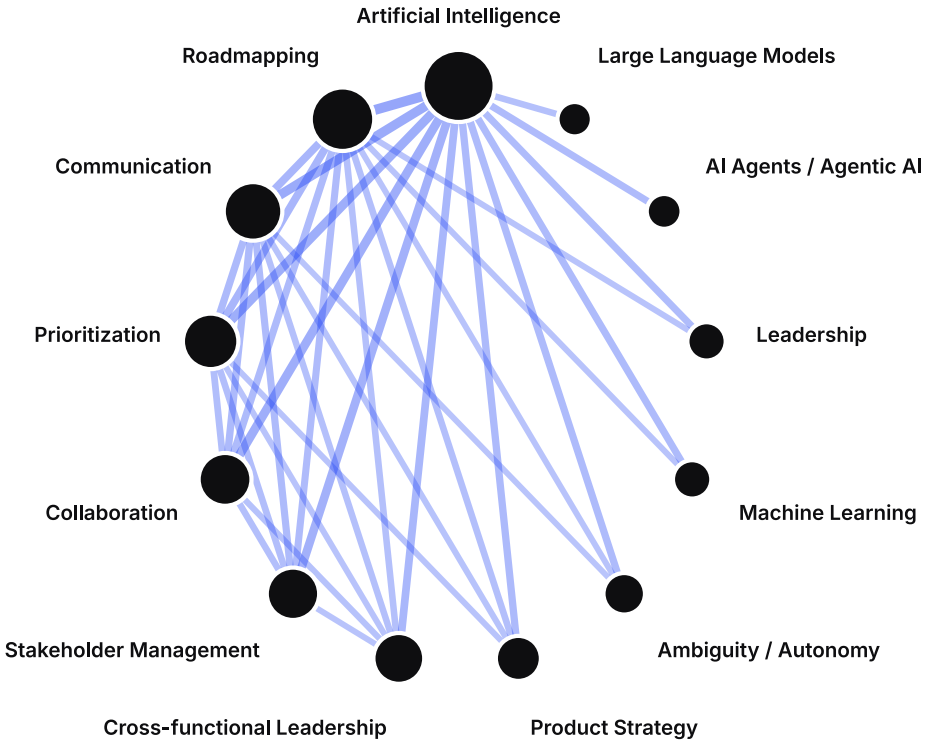

Of the five most-connected skills, 4 are product or human skills — AI capabilities attach to product craft rather than replacing it.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Pairs shown where both skills appear together in at least 15 postings.

Figure 22 — The AI Product Manager skill network, weighted by joint appearance.

SKILL A	SKILL B	POSTINGS WITH BOTH	LIFT
Artificial Intelligence	Roadmapping	1,849	1.00
Artificial Intelligence	Communication	1,723	1.00
Artificial Intelligence	Prioritization	1,641	1.00
Artificial Intelligence	Collaboration	1,621	1.00
Artificial Intelligence	Stakeholder Management	1,548	1.00
Roadmapping	Prioritization	1,401	1.12
Roadmapping	Communication	1,395	1.06
Artificial Intelligence	Cross-functional Leadership	1,381	1.00
Communication	Collaboration	1,290	1.12

SKILL A	SKILL B	POSTINGS WITH BOTH	LIFT
Artificial Intelligence	Product Strategy	1,289	1.00
Communication	Prioritization	1,287	1.10
Artificial Intelligence	Ambiguity / Autonomy	1,281	1.00
Roadmapping	Collaboration	1,279	1.03
Artificial Intelligence	Machine Learning	1,261	1.00

Table 13 — Lift above 1.0 means the pair appears together more often than independent frequency would predict.

Regional differences

Requirements are more similar across borders than candidates assume

Percentage of postings in each market that request the skill

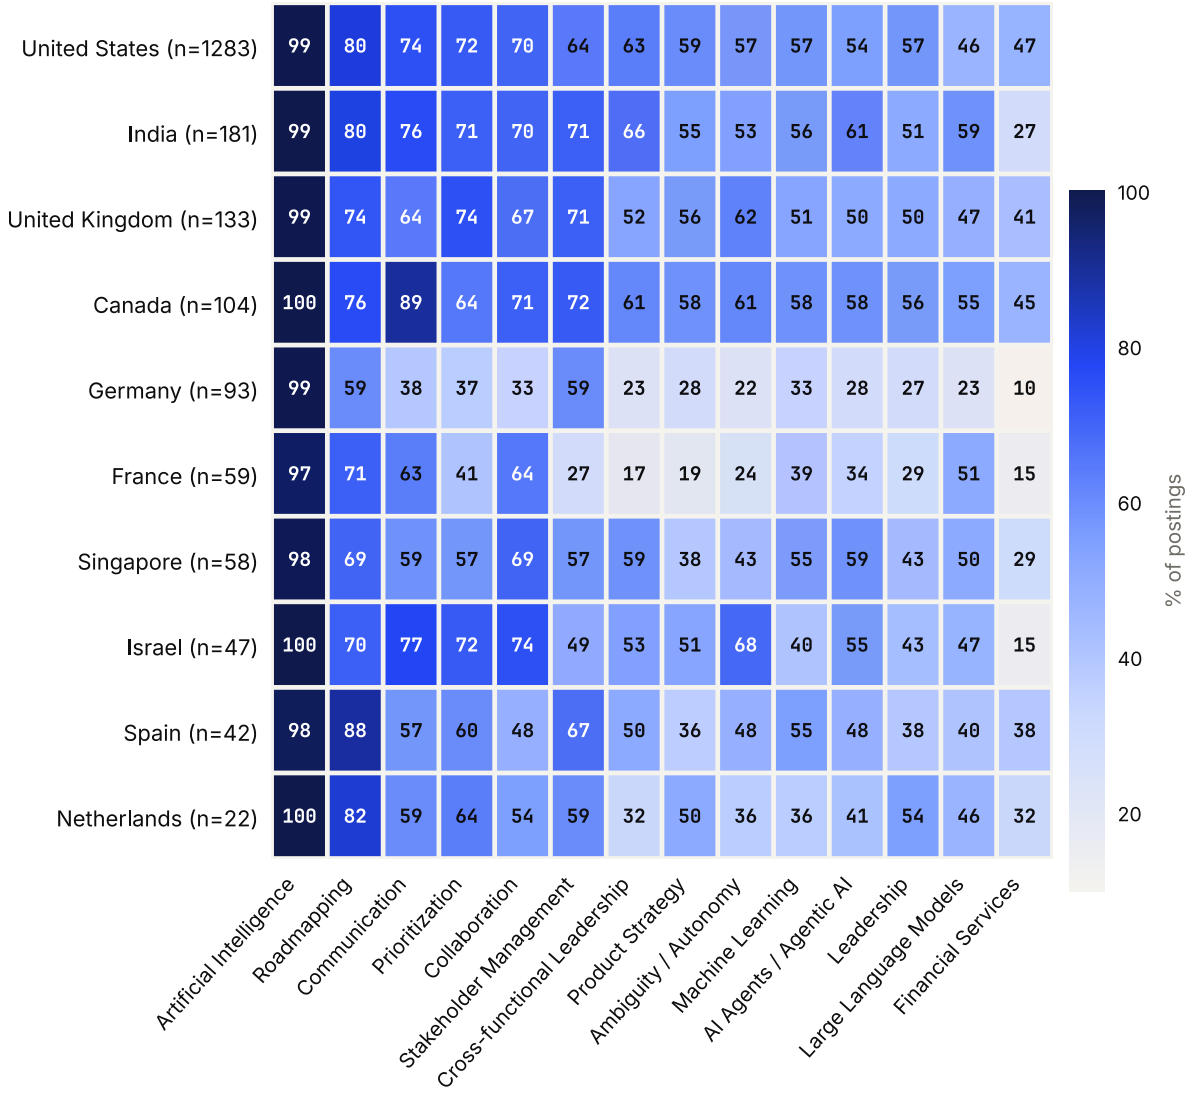

Requirements are largely standardised across markets — Artificial Intelligence varies by only 1pp between countries, while Cross-functional Leadership varies by 17pp.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Markets with at least 20 postings.

Figure 23 — Skill request rates across the largest hiring markets.

WHAT THE DATA SAYS

The variation between markets is narrower than the variation between seniority levels inside a single market. A senior posting in one country resembles a senior posting in another far more than it resembles a mid-level posting next door.

SECTION 08

08 The modern AI stack

LLMs, RAG, agents, evaluations, fine-tuning, vector search, MLOps and responsible AI — measured rather than assumed.

LLMs named in 46% of postings

RAG in 19.0%, agents in 51.2%

Frontier techniques remain minority requirements

This section exists because the gap between what candidates study and what employers write down is widest here.

The modern GenAI stack, measured

Share of unique 2026 AI PM postings referencing each AI or generative-AI capability

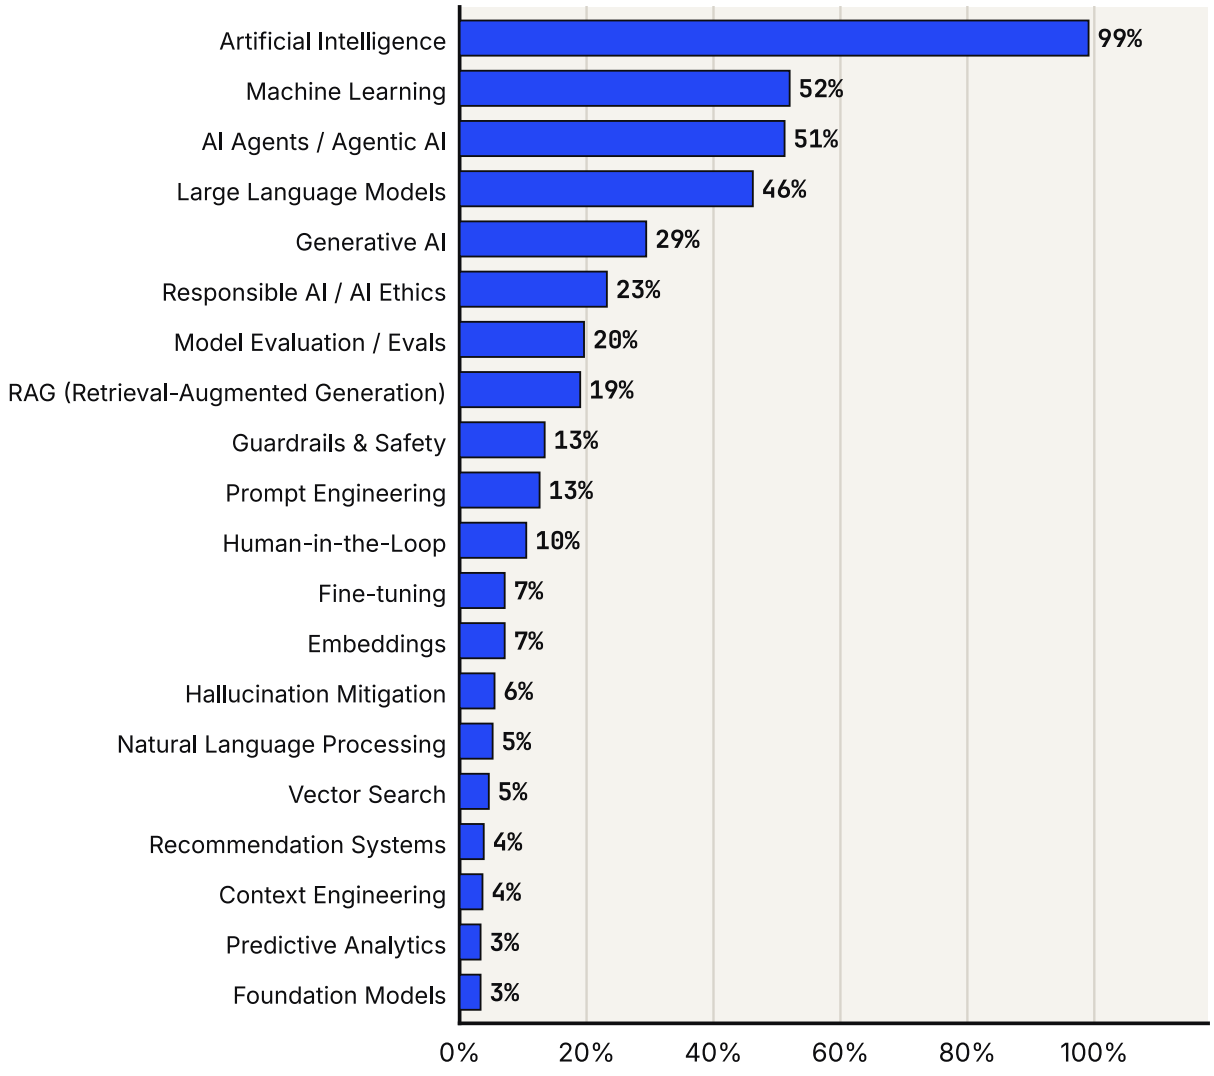

Artificial Intelligence appears in 99% of postings; the specialised GenAI techniques below it are still emerging requirements.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 24 — AI and generative-AI capabilities by share of postings.

RAG, agents and evaluations: how often are they really requested?

Share of unique 2026 AI PM postings naming each modern GenAI technique

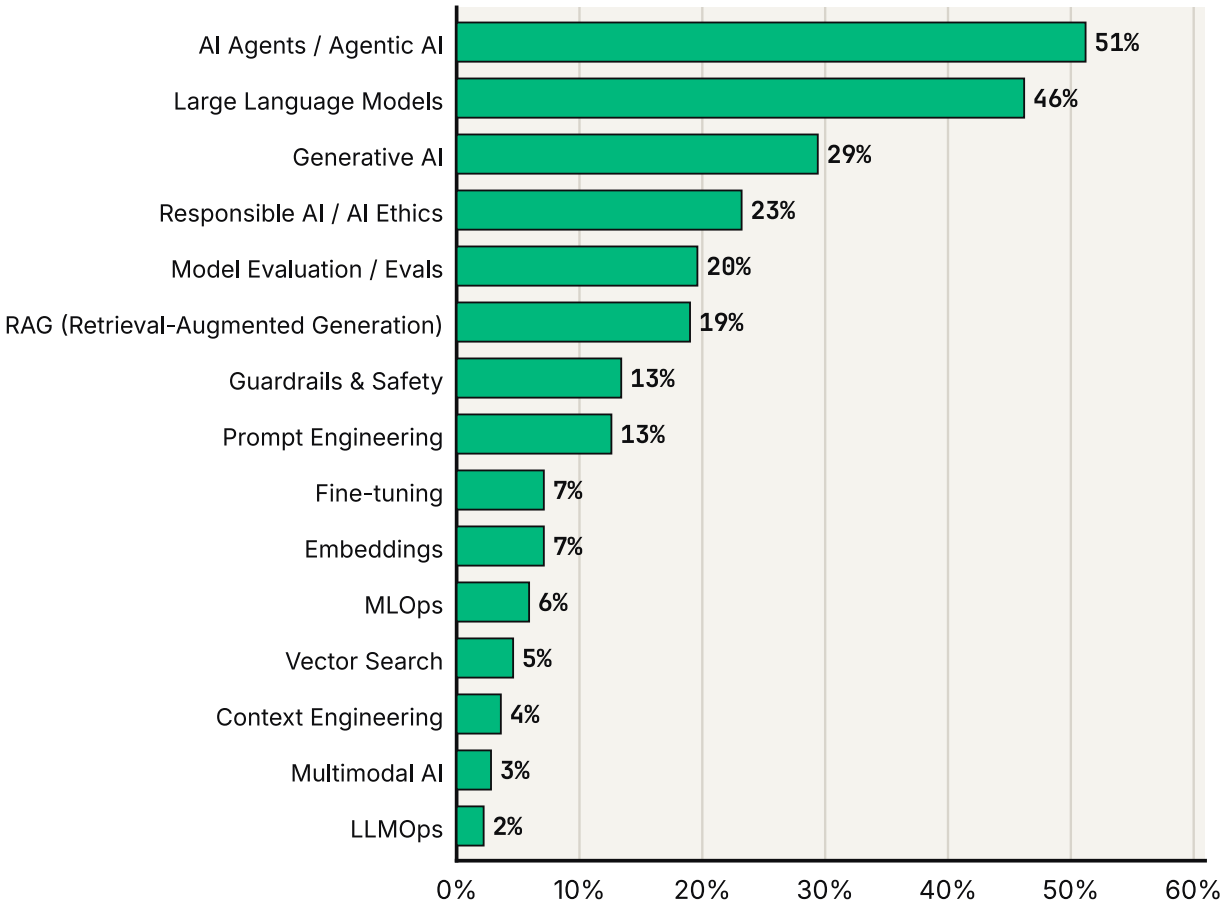

AI Agents / Agentic AI reaches 51% of postings, but only 4 of these 15 techniques clear 20% — the frontier stack is still a differentiator, not a baseline.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Measured on normalised terms including abbreviations (e.g. RAG, retrieval-augmented generation).

Figure 25 — How often modern GenAI techniques are actually named.

CAPABILITY	% OF POSTINGS	HOW IT SHOWS UP IN PRACTICE
Large Language Models	46.2%	The baseline vocabulary of the field. If a posting mentions any modern AI concept, it usually mentions this one.
Generative AI	29.4%	Frequently used as an umbrella term rather than a technical requirement; treat its presence as a signal about the produc
AI Agents / Agentic AI	51.2%	The fastest-moving category in the dataset. Where it appears it is usually accompanied by tool-use, permissioning and fa
RAG (Retrieval-Augmented Generation)	19.0%	Appears far less often than candidate discourse implies, but almost always in roles with real retrieval surfaces — searc

CAPABILITY	% OF POSTINGS	HOW IT SHOWS UP IN PRACTICE
Prompt Engineering	12.6%	Rarely a standalone requirement in product roles. It appears as one component of a broader design responsibility.
Model Evaluation / Evals	19.6%	The most under-rated item in the table. Where it appears it is close to non-negotiable, because it is how the employer d
Fine-tuning	7.1%	A specialist requirement concentrated in platform and applied-science-adjacent roles.
Vector Search	4.6%	Infrastructure-layer language; usually implies the PM owns retrieval quality.
Embeddings	7.1%	Almost always co-occurs with retrieval or search responsibilities.
Responsible AI / AI Ethics	23.2%	Concentrated in regulated industries and larger employers, where it is often non-negotiable.
Guardrails & Safety	13.4%	Appears as a product requirement, not a research topic: what the system must refuse to do.
Context Engineering	3.6%	Newest term in the ontology and the smallest count; included to establish a 2026 baseline.

Table 14 — Share of Tier A + Tier B postings naming each capability, normalised across synonyms and abbreviations.

FOR JOB SEEKERS

Learn evaluation before you learn fine-tuning. Evaluations appear in 19.6% of postings, transfer across every archetype and every industry, and are the one AI skill that is almost never listed as a nice-to-have. Fine-tuning (7.1%) is narrower, more expensive to learn, and concentrated in a small set of employers.

Responsible AI and governance

Responsible AI appears in 23.2% of postings and guardrails in 13.4%. Both are concentrated in regulated sectors. Treat these as a sector signal: candidates targeting healthcare, financial services or public-sector employers should be able to discuss model risk, auditability and human oversight in product terms.

SECTION 09

09 How technical do AI Product Managers need to be?

Python, SQL, APIs, cloud and architecture — what is genuinely requested, and what candidates over-prepare for.

84% of postings name no programming language

Python 9%, SQL 11%, APIs 28%

4% ask for Python and SQL together

The honest answer is: technical enough to ask the right question, almost never technical enough to write the production code.

How technical does an AI Product Manager have to be?

Share of unique 2026 AI PM postings by technical-requirement pattern

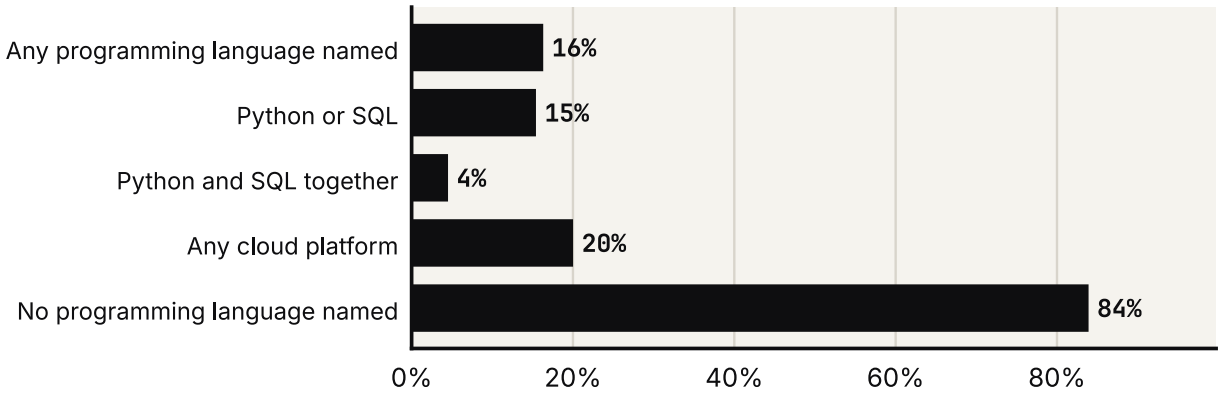

84% of AI PM postings name no programming language at all — fluency beats fluency in code.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. A mention means the posting names the technology; it does not imply production-level coding.

Figure 26 — Technical requirement patterns across AI Product Manager postings.

84% of postings name no programming language at all. 15% name Python or SQL; 4% name both. 20% mention a cloud platform or cloud concepts.

Which programming languages AI PM postings actually name

Share of unique 2026 postings mentioning each language anywhere in the role description

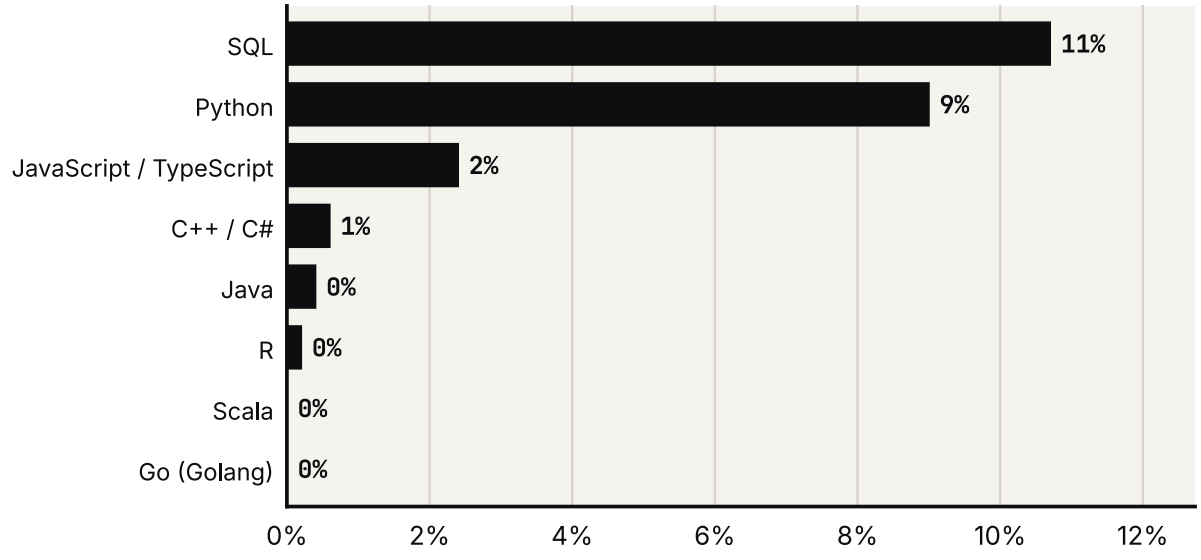

Python (9%) and SQL (11%) are the only languages that appear at scale.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. A mention indicates the language appears in the posting; it does not imply production coding is expected.

Figure 27 — Programming languages named in AI Product Manager postings.

Cloud and infrastructure requirements

Share of unique 2026 AI PM postings mentioning each cloud or platform capability

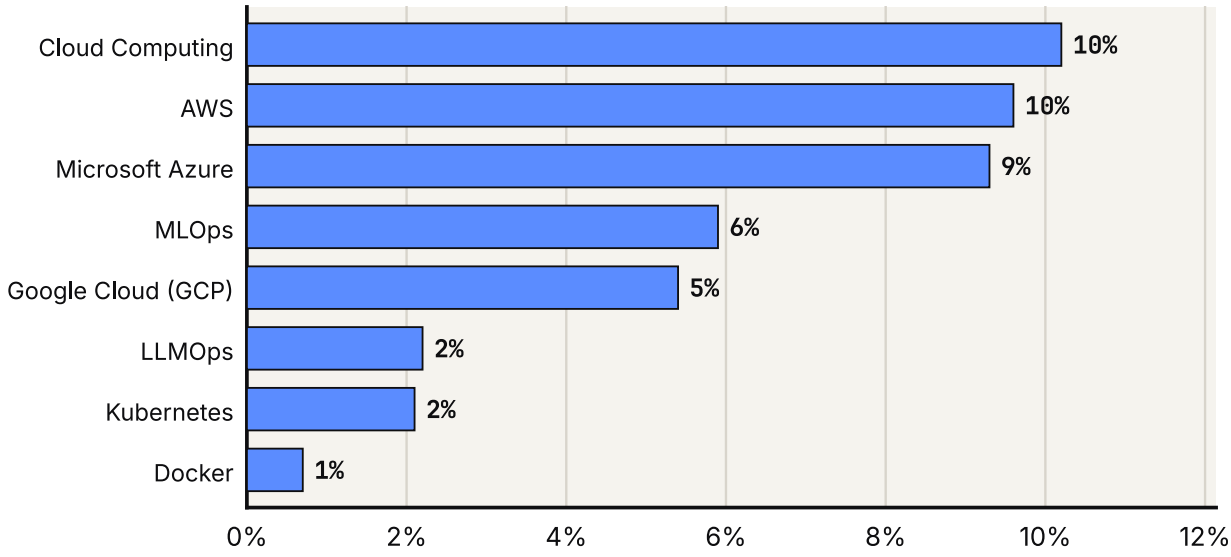

Generic cloud literacy (10%) is asked for far more often than any single vendor.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026.

Figure 28 — Cloud, container and operations capabilities named in postings.

Reading the technical signals correctly

1. **A mention is not a bar.** Postings name Python far more often than they describe a coding task. In most descriptions the surrounding language is analytical — query the data, prototype with an API, read a notebook — not engineering.
2. **SQL is the most honest technical requirement.** Where it appears it is usually tied to a concrete responsibility: own the metrics, interrogate the funnel, validate the model's impact yourself.
3. **APIs are the real technical literacy test.** Understanding rate limits, latency, token cost, retries and failure modes matters more in these postings than language syntax.
4. **Architecture language scales with seniority.** System design appears far more in principal and platform roles than in mid-level ones.

TECHNICAL CAPABILITY	% OF POSTINGS	REQUIRED SHARE
Python	9.0%	80%
SQL	10.7%	81%
APIs	28.2%	80%
REST APIs	28.3%	63%

TECHNICAL CAPABILITY	% OF POSTINGS	REQUIRED SHARE
Cloud Computing	10.2%	68%
AWS	9.6%	66%
Microsoft Azure	9.3%	69%
Google Cloud (GCP)	5.4%	65%
System Design / Architecture	10.9%	84%
Data Pipelines / ETL	13.2%	66%
MLOps	5.9%	79%
Docker	0.7%	60%
Kubernetes	2.1%	62%
Git / Version Control	3.6%	76%
Big Data (Spark/Hadoop)	6.3%	63%

Table 15 — Technical capability request rates.

WHAT THIS MEANS

The practical bar is: you can read an API reference, design an evaluation in a spreadsheet, query your own product data, and hold a credible conversation about latency and cost. That is achievable in months, not years — and it is what the postings actually describe.

$193K

median posted salary, USD-normalised, across 904 postings that disclose a range

Posted pay is high and widely dispersed: the 90th percentile reaches $265K. Every figure in this part is what an employer advertised, not what a candidate accepted.

PART III

The Money

What these roles pay, where they pay it, and what separates the top of the market from the middle.

10 Salary report

11 The premium segment

12 Industries and employers

13 Working models

14 Geography

15 High-paying roles

SECTION 10

10 Salary report

What AI Product Manager roles pay, measured from salaries disclosed in the postings themselves, reported in local currency and USD-normalised.

37% of postings disclose a pay range

Global median posted midpoint: $193K

Observed postings separated from external benchmarks throughout

Across 904 postings that disclose a range (37% of the sample), the global median posted midpoint is $193K USD-normalised.

Two rules govern every number in this section. First, these are **posted** salaries — what employers advertised, not what candidates accepted, and base pay rather than total compensation unless the posting said otherwise. Second, cross-country comparisons are nominal USD conversions at spot rates; they are not adjusted for cost of living, tax, or equity practice, and should not be read as statements about living standards.

$159K

25th percentile

USD-normalised midpoint

$193K

Median

n=904

$226K

75th percentile

USD-normalised midpoint

$265K

90th percentile

USD-normalised midpoint

What AI Product Managers are actually offered

Distribution of posted salary midpoints, USD-normalised (n=904 postings disclosing pay)

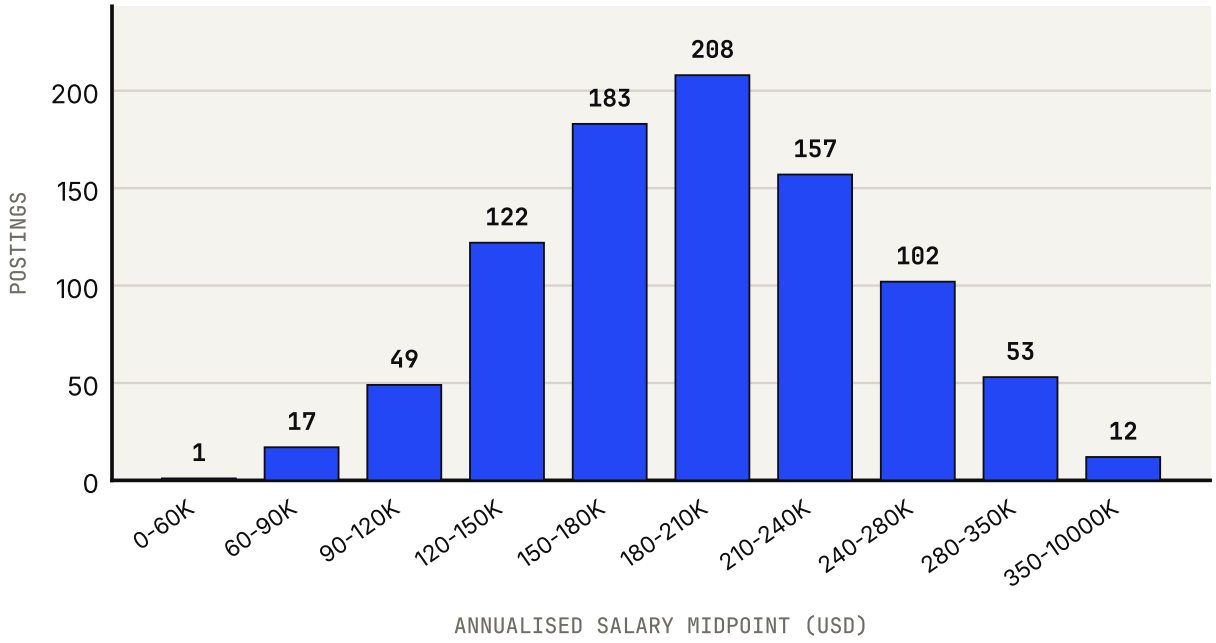

Global median posted midpoint is $193K; the 90th percentile reaches $265K.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Observed posting salaries only. Coverage: 37% of the sample. Converted at spot rates on 12 Sep 2026.

Figure 29 — Distribution of posted salary midpoints, USD-normalised.

Pay by seniority

Posted pay rises sharply with seniority

USD-normalised posted salary midpoint by seniority level; box shows interquartile range, line shows median

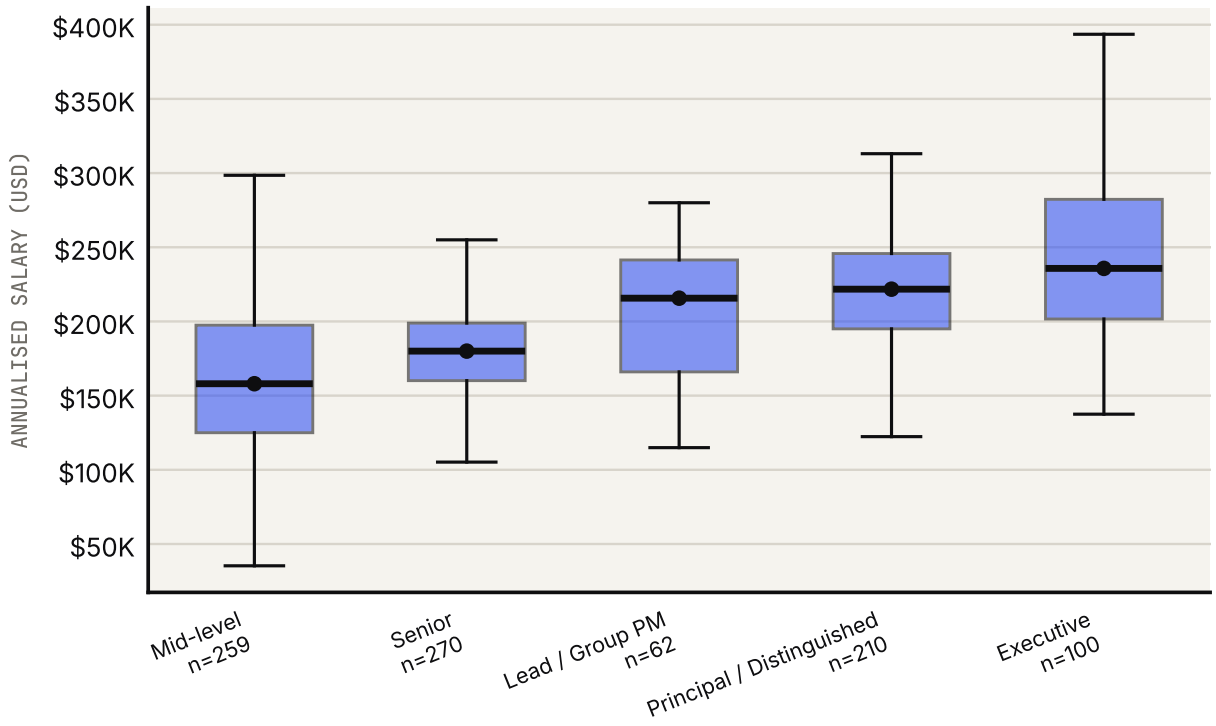

The largest single step is Senior to Lead / Group PM: $180K to $216K, a 20% jump in median posted pay.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Outliers suppressed. Levels with fewer than 8 disclosing postings omitted.

Figure 30 — Posted salary midpoint by seniority level.

SENIORITY	POSTINGS	25TH PCT	MEDIAN	75TH PCT
Mid-level	258	$125,000	$159,010	$197,500
Senior	270	$160,100	$180,000	$198,875
Lead / Group PM	62	$166,000	$215,679	$241,438
Principal / Distinguished	210	$195,000	$221,750	$245,762
Executive (VP/Head/Director)	100	$201,650	$235,750	$282,300

Table 16 — USD-normalised posted salary midpoints by seniority.

Pay by country

Where AI Product Managers are paid the most

Median USD-normalised posted salary midpoint by country

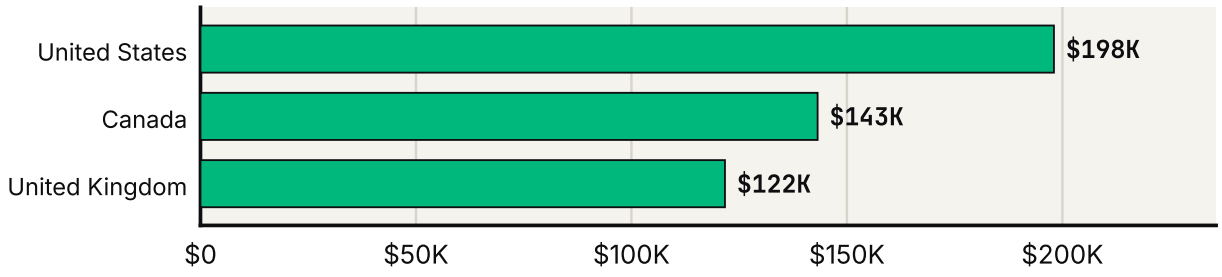

United States leads at $198K; nominal gaps overstate real differences in living standards.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Local-currency medians and sample sizes are reported in the Appendix. Cross-country comparison is nominal USD, not cost-of-living adjusted.

Figure 31 — Median posted salary by country, USD-normalised.

COUNTRY	N	MEDIAN (USD)	75TH (USD)	90TH (USD)	LOCAL MEDIAN
United States	774	$198,038	$231,000	$270,000	USD 198,037
Canada	63	$143,200	$168,615	$205,540	USD 155,000
United Kingdom	13	$121,701	$163,080	$210,529	GBP 95,500

Table 17 — Markets with at least 12 pay-disclosing postings. Local median is the median annualised midpoint in the market's modal posting currency.

READ THIS CAREFULLY

Nominal USD conversion overstates real differences. A posting in India at a lower USD figure may represent a stronger local position than a mid-range United States posting. Compare within a market first, across markets second.

Pay and experience

Experience sets a floor on pay, not a ceiling

Posted salary midpoint against minimum required experience (n=704 postings with both fields)

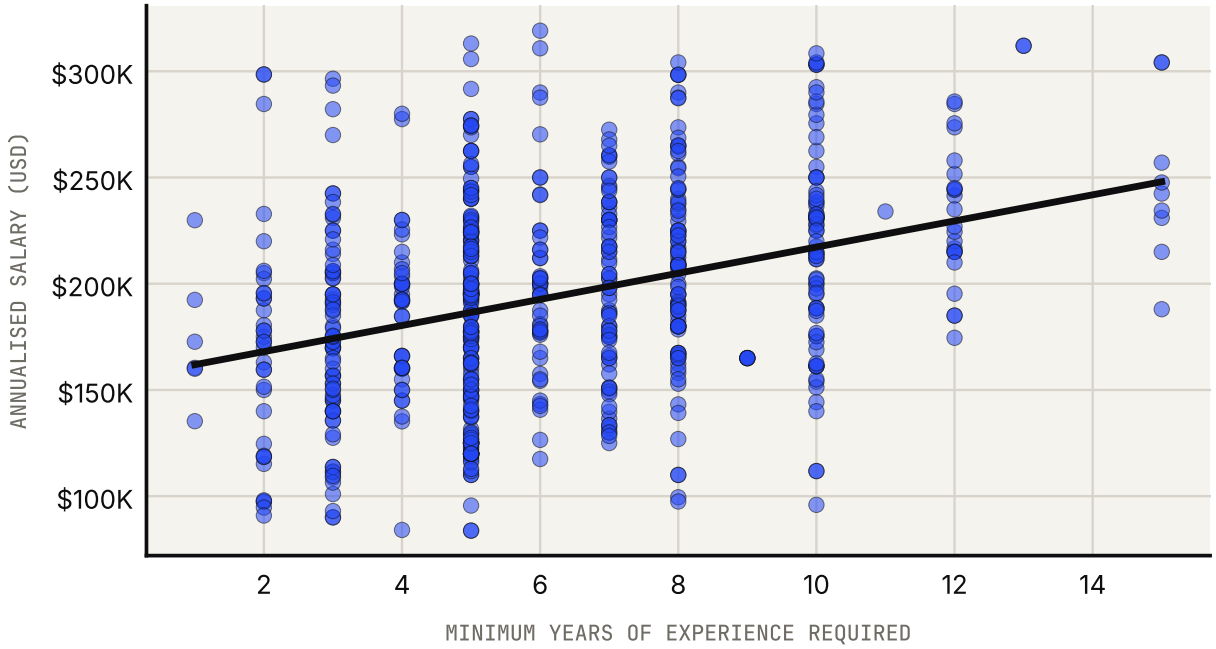

Correlation is real but partial (r = 0.11): experience sets the floor, scope and skill mix set the ceiling.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Each point is one posting. Line is an ordinary least-squares fit. The top 2% and bottom 1% of posted values are omitted so the trend stays legible.

Figure 32 — Posted salary against minimum required experience.

MINIMUM YEARS REQUIRED	POSTINGS	MEDIAN POSTED SALARY (USD)
2	37	$172,650
3	77	$172,650
4	47	$184,700
5	201	$185,000
6	50	$198,950
7	73	$198,038
8	91	$208,800
10	79	$225,000

MINIMUM YEARS REQUIRED	POSTINGS	MEDIAN POSTED SALARY (USD)
12	23	$226,966
15	12	$252,325

Table 18 — Bands with at least 8 pay-disclosing postings.

Pay by industry, archetype and working model

Posted pay by industry

Median USD-normalised posted salary midpoint among pay-disclosing postings

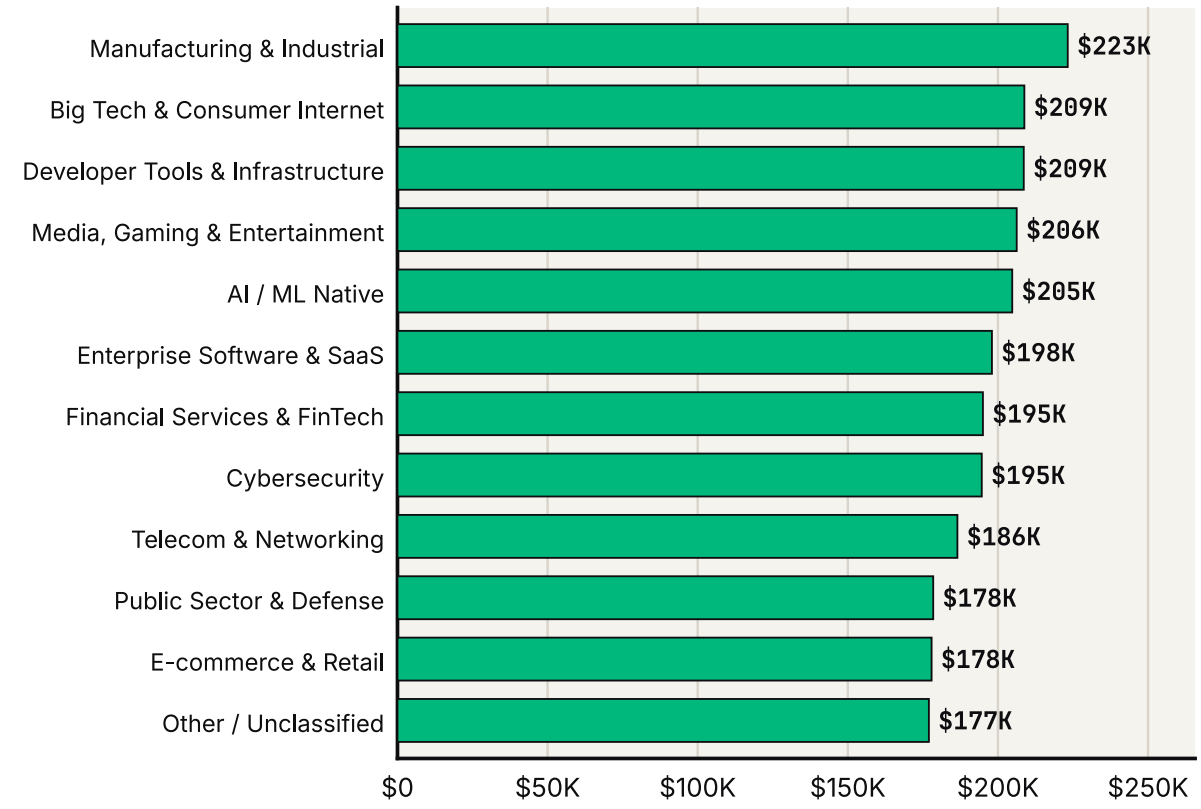

Manufacturing & Industrial pays the highest median at $223K.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Industries with at least 10 pay-disclosing postings.

Figure 33 — Median posted salary by employer industry.

INDUSTRY	POSTINGS	MEDIAN POSTED SALARY (USD)
Manufacturing & Industrial	22	$223,250
Big Tech & Consumer Internet	27	$208,800
Developer Tools & Infrastructure	74	$208,600

INDUSTRY	POSTINGS	MEDIAN POSTED SALARY (USD)
Media, Gaming & Entertainment	20	$206,250
AI / ML Native	34	$204,762
Enterprise Software & SaaS	200	$198,000
Financial Services & FinTech	134	$195,000
Cybersecurity	36	$194,612
Telecom & Networking	15	$186,500
Public Sector & Defense	14	$178,440
E-commerce & Retail	49	$177,900
Other / Unclassified	115	$177,000
Consulting & Professional Services	14	$172,650
Healthcare & Life Sciences	105	$162,850

Table 19 — Industries with at least 10 pay-disclosing postings.

ROLE ARCHETYPE	POSTINGS	MEDIAN POSTED SALARY (USD)
AI Platform / Infra PM	86	$204,275
Agentic AI PM	328	$195,000
Data / ML Product PM	119	$195,000
General AI Product Manager	44	$193,971
Conversational AI PM	31	$187,217
GenAI / LLM Product PM	216	$185,000
Responsible AI / Governance PM	17	$177,900
Applied AI / AI Applications PM	63	$175,300

Table 20 — Archetypes with at least 8 pay-disclosing postings.

WORKING MODEL	POSTINGS	MEDIAN POSTED SALARY (USD)
Hybrid	237	$198,500
Not specified	341	$191,250
On-site	164	$196,388
Remote	162	$180,000

Table 21 — Posted salary by stated working model.

Which skills sit alongside higher pay

The skills associated with higher posted pay

Difference in median posted salary between postings that name the skill and those that do not — US postings

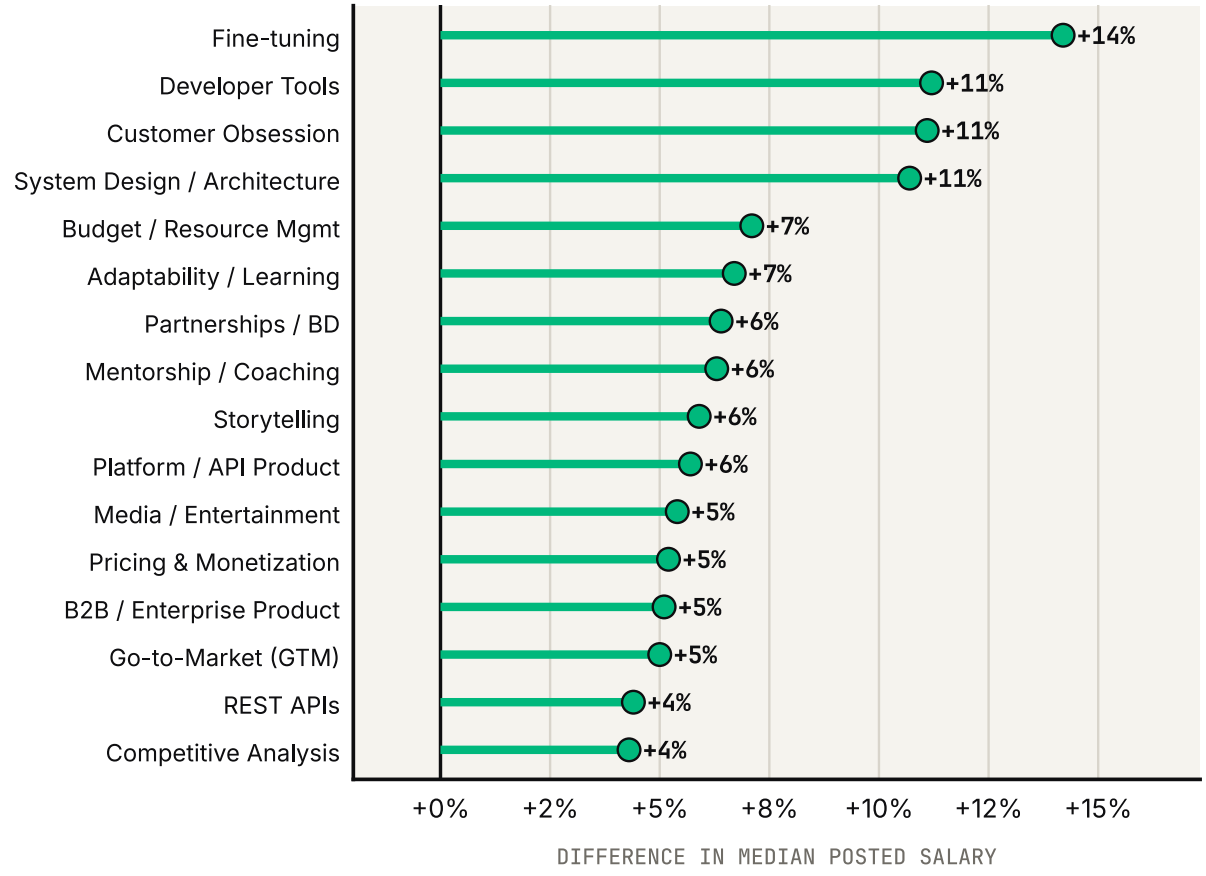

Fine-tuning shows the largest gap (+14%, $223K vs $196K) — association, not causation: scope and seniority move with skills.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Association, not causation. Basis pool n=774; skills shown where at least 15 postings on each side disclose pay.

Figure 34 — Difference in median posted salary between postings naming a skill and those that do not.

SKILL	N WITH SKILL	MEDIAN WITH	MEDIAN WITHOUT	DIFFERENCE
Fine-tuning	66	$223,225	$195,546	+14%
Developer Tools	60	$217,500	$195,650	+11%
Customer Obsession	91	$216,650	$195,000	+11%

SKILL	N WITH SKILL	MEDIAN WITH	MEDIAN WITHOUT	DIFFERENCE
System Design / Architecture	97	$216,650	$195,650	+11%
Budget / Resource Mgmt	64	$210,500	$196,500	+7%
Adaptability / Learning	132	$208,800	$195,650	+7%
Partnerships / BD	229	$207,500	$195,000	+6%
Mentorship / Coaching	194	$207,200	$195,000	+6%
Storytelling	120	$207,100	$195,650	+6%
Platform / API Product	144	$206,700	$195,546	+6%
Media / Entertainment	146	$206,250	$195,650	+5%
Pricing & Monetization	91	$205,900	$195,650	+5%
B2B / Enterprise Product	323	$205,000	$195,000	+5%
Go-to-Market (GTM)	264	$204,775	$195,000	+5%
REST APIs	275	$203,500	$195,000	+4%
Competitive Analysis	85	$205,900	$197,500	+4%

Table 22 — Computed on US postings to hold geography broadly constant. Association, not causation.

Controlling for geography and seniority

A log-linear regression on 901 pay-disclosing postings, controlling for country, seniority band, stated experience and working model, explains 51% of variance in posted pay (adjusted R^2 0.48). The table reports skill coefficients that remain statistically significant once those controls are applied.

FEATURE (SKILL)	ESTIMATED EFFECT ON POSTED PAY	P-VALUE
Roadmapping	+4.9%	0.046
REST APIs	+4.3%	0.014
Mentorship Coaching	+4.3%	0.028
Machine Learning	+3.4%	0.042
Financial Services	-3.4%	0.049
Communication	-4.4%	0.032
Metrics KPIs	-4.7%	0.007
Product Requirements PRDs	-4.8%	0.003
HR Future of Work	-6.4%	0.004

Table 23 — Coefficients from an OLS model on log posted-salary midpoint with heteroskedasticity-robust standard errors. Effects are associations within this sample, not causal estimates.

Observed postings versus external benchmarks

Every salary figure above is observed directly in a job posting. Where readers want a cross-check against self-reported compensation data — which includes equity and bonus and therefore runs higher — the established public sources are Levels.fyi,

Glassdoor, Built In and, for India, AmbitionBox. Those are different measurements of a different quantity and are not merged into this dataset. The Appendix explains why.

SECTION 11

11 The $200K+ AI Product Manager

What separates the highest-paying segment from everything else — experience, scope, skills, industries and employers.

44% of pay-disclosing postings sit at $200K or above

Scope language, not tooling, is the dividing line

Analysed against all lower-paying postings in the same sample

393 postings in this sample advertise a midpoint at $200K or above — 44% of everything that discloses pay. Their median is $232K.

393

postings at $200K+

44% of pay-disclosing

$232K

median in this segment

USD-normalised

7 yrs

median experience asked

vs 5 years elsewhere

20%

state a degree as required

MBA named in 12.0%

What the $200K+ roles ask for that other roles do not

Percentage-point difference in skill frequency: postings at $200K+ (n=393) versus all lower-paying postings

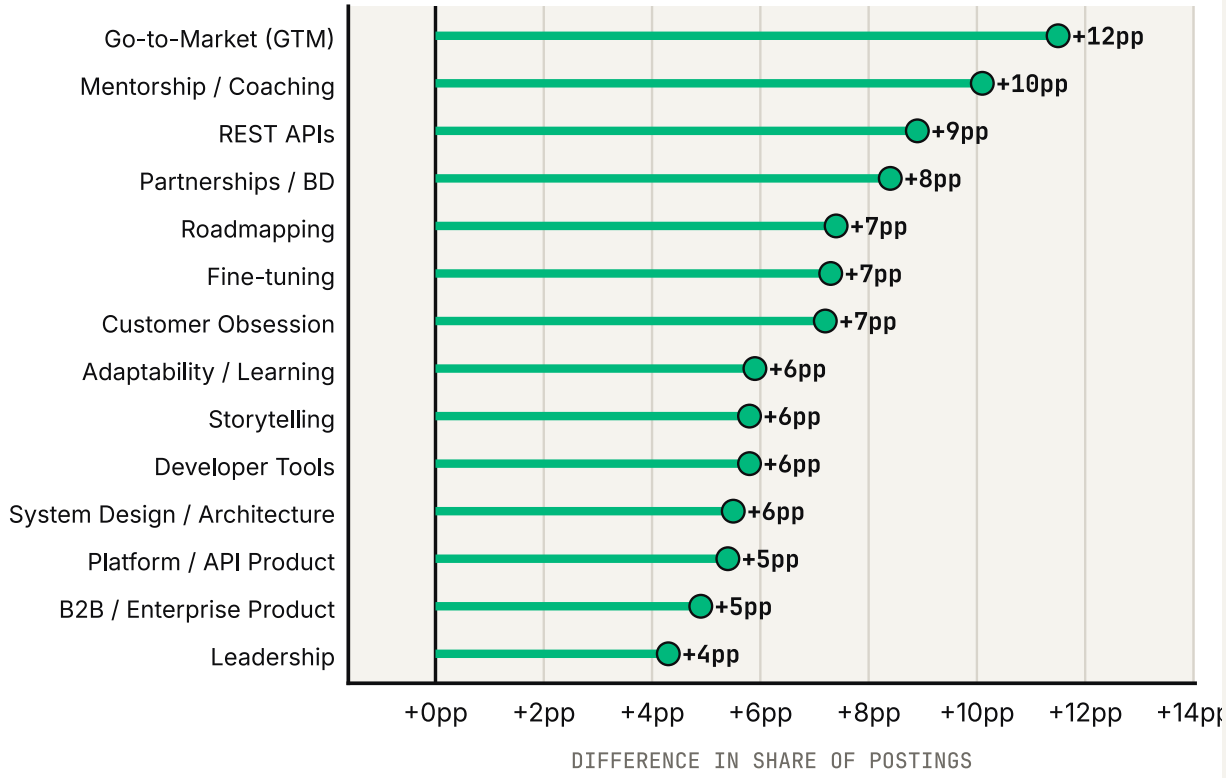

Go-to-Market (GTM) appears in 40% of premium roles versus 28% elsewhere; median required experience is 7 years versus 5.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. 44% of pay-disclosing postings sit at or above $200K. Median in this segment: $232K.

Figure 35 — Skills over-represented in $200K+ postings relative to all lower-paying postings.

What actually differs

SKILL	$200K+ POSTINGS	ALL OTHERS	DIFFERENCE
Go-to-Market (GTM)	39.7%	28.2%	+11.5pp
Mentorship / Coaching	30.3%	20.2%	+10.1pp
REST APIs	40.2%	31.3%	+8.9pp
Partnerships / BD	33.8%	25.4%	+8.4pp
Roadmapping	83.7%	76.3%	+7.4pp
Fine-tuning	12.2%	4.9%	+7.3pp

SKILL	$200K+ POSTINGS	ALL OTHERS	DIFFERENCE
Customer Obsession	16.0%	8.8%	+7.2pp
Adaptability / Learning	20.6%	14.7%	+5.9pp
Storytelling	18.3%	12.5%	+5.8pp
Developer Tools	10.9%	5.1%	+5.8pp
System Design / Architecture	15.5%	10.0%	+5.5pp
Platform / API Product	22.6%	17.2%	+5.4pp
B2B / Enterprise Product	47.8%	42.9%	+4.9pp
Leadership	58.3%	54.0%	+4.3pp

Table 24 — Percentage-point difference in request frequency between the premium segment and all lower-paying postings.

Seniority composition

SENIORITY	POSTINGS	SHARE OF $200K+ SEGMENT
Principal / Distinguished	152	38.7%
Executive (VP/Head/Director)	77	19.6%
Senior	65	16.5%
Mid-level	62	15.8%
Lead / Group PM	37	9.4%
Entry / Associate	0	0.0%

Table 25 — Seniority mix within the premium segment.

INDUSTRY	POSTINGS	SHARE OF $200K+ SEGMENT
Enterprise Software & SaaS	99	25.2%
Financial Services & FinTech	63	16.0%
Developer Tools & Infrastructure	41	10.4%
Other / Unclassified	36	9.2%
Healthcare & Life Sciences	32	8.1%
AI / ML Native	21	5.3%
Big Tech & Consumer Internet	19	4.8%
Manufacturing & Industrial	16	4.1%
Cybersecurity	16	4.1%
Media, Gaming & Entertainment	14	3.6%

Table 26 — Industries represented in the premium segment.

Geographically the segment is concentrated: San Francisco, San Jose, New York, Seattle and San Diego account for the largest share of premium postings with a named city.

THE $200K+ DIFFERENCE

The dividing line is scope, not stack. The premium segment asks for 7 years of experience against 5 years elsewhere, and over-indexes on platform ownership, organisational leadership and systems-level responsibility. Adding another named framework to a CV does not move a candidate into this band. Owning a surface that other teams depend on does.

SECTION 12

12 Top industries and employers

Which sectors are buying AI product management, which companies post the most, and how AI-native employers differ from everyone else.

Enterprise Software & SaaS leads at 18% of postings

5% of postings show AI-native company language

Demand is spread across more than a dozen distinct sectors

The most useful thing this dataset says about employers is how many of them there are.

1,560 distinct organisations posted at least one AI Product Manager role. 76% posted exactly one. The concentration that candidates assume — a handful of labs absorbing the market — is not present in the data.

INDUSTRY	POSTINGS	SHARE OF SAMPLE
Enterprise Software & SaaS	450	18.4%
Financial Services & FinTech	417	17.0%
Healthcare & Life Sciences	197	8.0%
Developer Tools & Infrastructure	149	6.1%
E-commerce & Retail	97	4.0%
Big Tech & Consumer Internet	97	4.0%
Cybersecurity	91	3.7%
Manufacturing & Industrial	72	2.9%
AI / ML Native	60	2.5%
Consulting & Professional Services	52	2.1%

INDUSTRY	POSTINGS	SHARE OF SAMPLE
Media, Gaming & Entertainment	51	2.1%
Public Sector & Defense	37	1.5%
Telecom & Networking	33	1.3%
Automotive & Mobility	32	1.3%
Travel & Hospitality	27	1.1%
Real Estate & PropTech	24	1.0%

Table 27 — Industry assigned by weighted classification of posting content, reconciled to one label per employer.

Sector notes

Enterprise Software & SaaS

18.4% of postings

Fastest to adopt AI PM titles. Screens hardest on product fundamentals and go-to-market.

Financial Services & FinTech

17.0% of postings

Governance, model risk and auditability appear far above sample average. Domain fluency is a genuine filter.

Healthcare & Life Sciences

8.0% of postings

Regulatory literacy and human-in-the-loop design appear consistently; clinical workflow knowledge is often decisive.

Developer Tools & Infrastructure

6.1% of postings

The most technical AI PM postings in the dataset. API and system-design language runs well above average.

E-commerce & Retail

4.0% of postings

Ranking, personalisation and merchandising; classical ML remains central alongside generative work.

Big Tech & Consumer Internet

4.0% of postings

Highest bar on scale, experimentation and metrics rigour; also the most competitive applicant pool.

Cybersecurity

3.7% of postings

Detection quality, false-positive economics and evaluation appear more than generative language.

Manufacturing & Industrial 2.9% of postings

Operational and reliability framing; predictive maintenance and quality use cases dominate.

AI / ML Native 2.5% of postings

Fastest-moving briefs, least structured processes, highest expectation of independent judgement.

Consulting & Professional Services 2.1% of postings

Breadth over depth; client-facing communication weighted heavily.

Employers

EMPLOYER	POSTINGS	MEDIAN POSTED PAY (USD)	PRIMARY INDUSTRY
Amazon	32	$195,650	E-commerce & Retail
JPMorgan Chase	24	—	Financial Services & FinTech
Google	24	—	Big Tech & Consumer Internet
Airwallex	17	$230,000	Financial Services & FinTech
Veeva	16	$125,000	Healthcare & Life Sciences
Palo Alto Networks	16	$215,000	Cybersecurity
Microsoft	15	$208,800	Big Tech & Consumer Internet
Amazon Web Services (AWS)	14	$213,050	Enterprise Software & SaaS
OKX	13	—	Financial Services & FinTech
Capital One	13	$176,400	Financial Services & FinTech
Veeva Systems	13	$125,000	Healthcare & Life Sciences
TikTok	13	—	Big Tech & Consumer Internet
ServiceNow	12	$304,150	Enterprise Software & SaaS
Meta	12	—	Big Tech & Consumer Internet
Qualcomm	11	$245,000	Manufacturing & Industrial
Adobe	11	$238,450	Enterprise Software & SaaS
Intuit	11	$241,750	Financial Services & FinTech
BNY	11	—	Financial Services & FinTech
Citi	11	$228,000	Financial Services & FinTech
Walmart	11	$165,000	E-commerce & Retail

Table 28 — Median pay shown only where at least five of an employer's postings disclose a range.

SECTION 13

13 Remote, hybrid or office

How AI Product Manager roles are structured, how that varies by market and seniority, and what it does to pay.

Hybrid 25% · Remote 13% · On-site 15%

48% state no working model at all

Working model varies more by country than by seniority

13% of AI Product Manager postings are fully remote and 25% hybrid. A further 48% name a location without describing a working model — itself a finding about how these roles are advertised.

Remote, hybrid or in the office?

Share of unique 2026 AI PM postings by stated working model

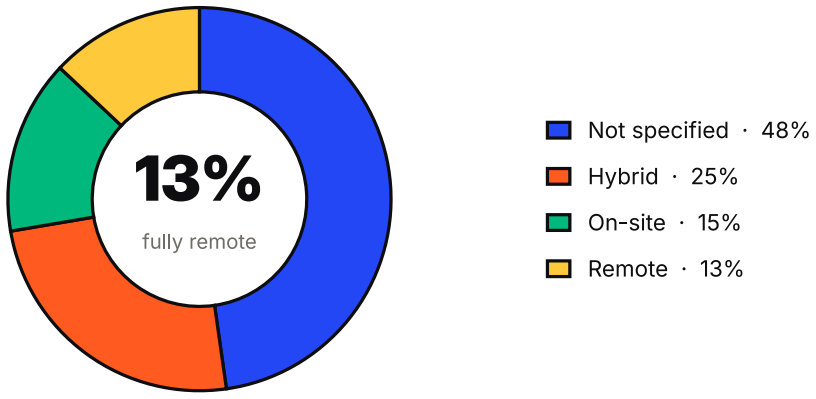

13% of postings are fully remote and 25% hybrid; 15% are on-site and 48% state no working model at all.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. 'Not specified' covers postings that name a location without describing a working model.

Figure 36 — Working models across AI Product Manager postings.

Working models differ sharply by country

Share of unique 2026 AI PM postings by working model within each market

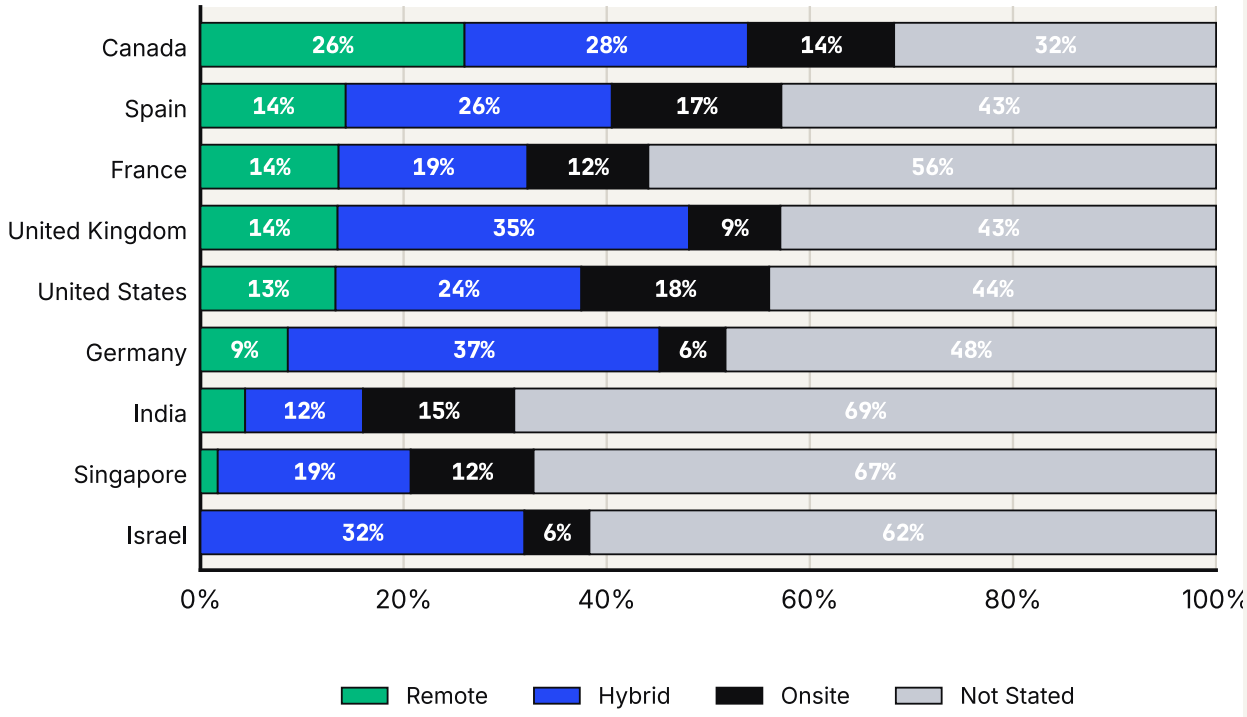

Remote share ranges from 0% (Israel) to 26% (Canada) — working model is a national-market characteristic, not a global norm.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Percentages are of all postings in that market, including those stating no working model.

Figure 37 — Working model composition within the largest hiring markets.

COUNTRY	POSTINGS	REMOTE	HYBRID	ON-SITE
Canada	104	26.0%	27.9%	14.4%
Spain	42	14.3%	26.2%	16.7%
France	59	13.6%	18.6%	11.9%
United Kingdom	133	13.5%	34.6%	9.0%
United States	1,283	13.3%	24.2%	18.5%
Germany	93	8.6%	36.6%	6.5%
India	181	4.4%	11.6%	14.9%
Singapore	58	1.7%	19.0%	12.1%
Israel	47	0.0%	31.9%	6.4%

Table 29 — Markets with at least 30 postings. Percentages are of all postings in the market, including those stating no model.

SENIORITY	POSTINGS	REMOTE	HYBRID	ON-SITE
Entry / Associate	37	8.1%	8.1%	16.2%
Mid-level	1,002	11.2%	23.1%	13.6%
Senior	648	15.0%	27.2%	13.0%
Lead / Group PM	191	8.4%	33.5%	20.4%
Principal / Distinguished	344	18.9%	24.7%	13.7%
Executive (VP/Head/Director)	226	11.1%	18.6%	21.2%

Table 30 — Working model by seniority band.

WHAT THIS MEANS

Remote availability is a property of the employer's market and operating model, not of the AI PM role itself. Candidates optimising for remote work should filter by country and company type before filtering by title.

SECTION 14

14 Geographic report

Market-by-market detail for the largest AI Product Manager economies: volume, seniority, experience, pay and working model.

53 countries represented

Requirements converge; seniority mix and pay diverge

Local currency reported alongside USD throughout

The skill requirements are close to globally standardised. What differs by market is the shape of demand: which seniority levels are being hired, how work is structured, and what it pays.

Requirements are more similar across borders than candidates assume

Percentage of postings in each market that request the skill

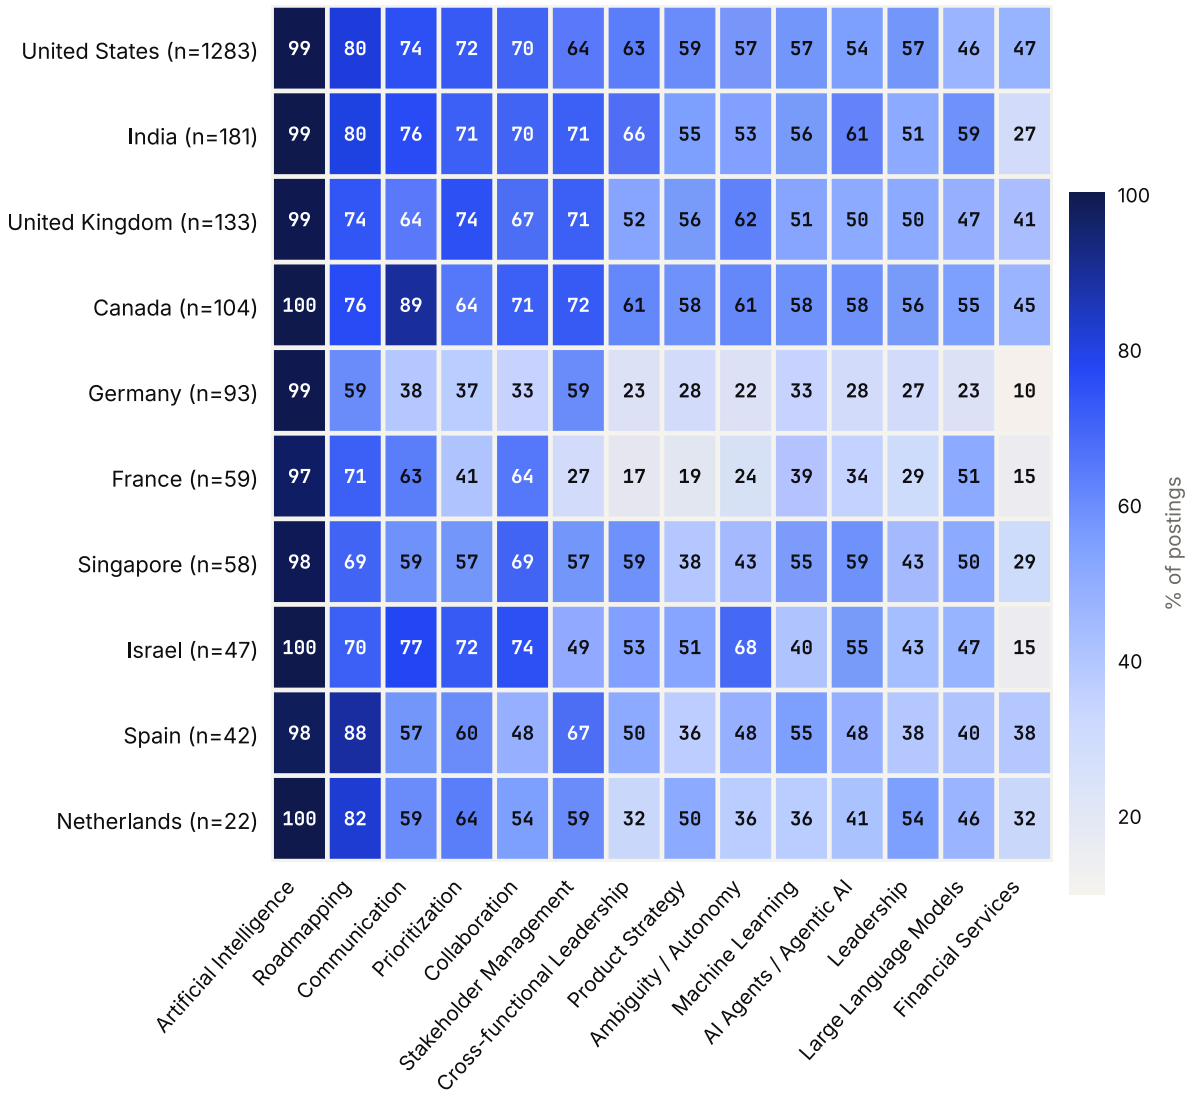

Requirements are largely standardised across markets — Artificial Intelligence varies by only 1pp between countries, while Cross-functional Leadership varies by 17pp.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Markets with at least 20 postings.

Figure 38 — Skill request rates across the largest hiring markets.

United States 1,283 postings

The deepest market by a wide margin and the most transparent on pay, largely because of state-level pay-disclosure laws. Highest concentration of platform and infrastructure roles.
1,283 postings (52.4% of global) · median posted pay $198K (n=774) · median experience 5 yrs · 13% remote

India — 181 postings

The fastest-scaling market in the sample and the most heterogeneous: global capability centres, domestic product companies and services firms all hiring under the same title with very different briefs.
181 postings (7.4% of global) · median experience 5 yrs · 4% remote

United Kingdom — 133 postings

Concentrated in London with a strong financial-services and enterprise-software mix; governance language appears above sample average.
133 postings (5.4% of global) · median posted pay $122K (n=13) · median experience 5 yrs · 14% remote

Canada — 104 postings

Toronto, Vancouver, Montreal and Waterloo; a high share of roles attached to United States-headquartered employers.
104 postings (4.2% of global) · median posted pay $143K (n=63) · median experience 5 yrs · 26% remote

Germany — 93 postings

Industrial, automotive and enterprise-software weighting; more on-site expectation than the sample average.
93 postings (3.8% of global) · median experience 5 yrs · 9% remote

France — 59 postings

Paris-centred, with a notable public and enterprise component.
59 postings (2.4% of global) · median experience 4 yrs · 14% remote

Singapore — 58 postings

The regional hub role: postings frequently carry APAC-wide scope.
58 postings (2.4% of global) · median experience 5 yrs · 2% remote

Israel — 47 postings

Disproportionately technical and product-platform oriented relative to its volume.
47 postings (1.9% of global) · median experience 5 yrs · 0% remote

Spain — 42 postings

Engineering-hub activity concentrated in Madrid and Barcelona; a high share of roles attached to non-Spanish employers.
42 postings (1.7% of global) · median experience 5 yrs · 14% remote

Netherlands — 22 postings

Small but disproportionately senior; English-language postings dominate.
22 postings (0.9% of global) · median experience 5 yrs

Mainland China — 20 postings

20 postings (0.8% of global) · median experience 5 yrs

Australia — 20 postings

Sydney and Melbourne; enterprise and financial services dominate.
20 postings (0.8% of global)

Poland — 20 postings

An engineering-centre market where product ownership, not just delivery, now appears in the postings.
20 postings (0.8% of global) · median experience 5 yrs

Portugal — 19 postings

19 postings (0.8% of global)

SECTION 15

15 What high-paying jobs ask for

A direct comparison between the top and bottom quartiles of posted pay, holding the sample and the extraction method constant.

Top quartile versus bottom quartile of posted pay

Differences measured in percentage points of request frequency

Read as association, not as a return on learning

The top quartile of pay-disclosing postings (n=227, median $257K) differs from the bottom quartile (n=226, median $129K) in ways that are consistent and mostly unglamorous.

SKILL	TOP QUARTILE	BOTTOM QUARTILE	DIFFERENCE
Mentorship / Coaching	35.2%	15.0%	+20.2pp
Roadmapping	84.6%	67.7%	+16.9pp
Leadership	61.7%	48.2%	+13.5pp
Go-to-Market (GTM)	36.6%	25.7%	+10.9pp
Partnerships / BD	36.1%	25.2%	+10.9pp
Customer Obsession	18.9%	8.8%	+10.1pp
Executive Communication	27.8%	18.1%	+9.7pp
Fine-tuning	15.0%	6.2%	+8.8pp
System Design / Architecture	16.7%	8.4%	+8.3pp
Cross-functional Leadership	64.8%	56.6%	+8.2pp
Critical Thinking	11.0%	3.1%	+7.9pp
Public Sector / Defense	25.1%	17.3%	+7.8pp
Adaptability / Learning	18.9%	11.1%	+7.8pp
Platform / API Product	22.5%	15.0%	+7.5pp
Product Strategy	59.5%	52.2%	+7.3pp
REST APIs	38.3%	32.3%	+6.0pp
Model Evaluation / Evals	23.3%	17.3%	+6.0pp
Machine Learning	57.7%	52.2%	+5.5pp

Table 31 — Skills most over-represented in the highest-paying quartile of postings.

What the bottom quartile asks for more

SKILL	TOP QUARTILE	BOTTOM QUARTILE	DIFFERENCE
Data Analysis	7.9%	20.4%	-12.5pp
Customer Empathy / UX	30.4%	43.4%	-13.0pp
Communication	69.2%	82.3%	-13.1pp
APIs	23.3%	37.6%	-14.3pp
Stakeholder Management	56.8%	72.6%	-15.8pp
Product Requirements (PRDs)	20.7%	36.7%	-16.0pp
Agile / Scrum	17.6%	35.0%	-17.4pp
Healthcare / Life Sciences	40.5%	58.4%	-17.9pp

Table 32 — Skills over-represented in the lowest-paying quartile.

WHAT THIS MEANS

The pattern across both tables is the same one that appears in the $200K+ analysis: what distinguishes better-paid postings is the size of the problem, not the novelty of the technology. Ownership language, platform scope, cross-organisational leadership and commercial accountability separate the quartiles more cleanly than any individual AI technique.

Employers are not hiring researchers. They are hiring product managers who are not afraid of models.

PART IV

The Signal

Underneath the frequency tables: which requirements are moving, how the role has split into archetypes, and what employers are really buying.

16 Momentum

17 Archetypes

18 What employers really want

SECTION 16

16 The skills gaining momentum

Like-for-like comparison between the 2025 and 2026 cohorts inside this dataset, using one extraction pipeline and one ontology.

Same dataset, same pipeline, two cohorts

No cross-source growth rates are reported

Percentage-point change, not relative growth, is the headline measure

The 2025 cohort in this dataset was too small to support a like-for-like trend comparison at the confidence level this report requires. Rather than assemble a growth rate from incompatible sources, this section reports the 2026 baseline only.

Section 8 establishes the 2026 baseline for every modern AI capability in the ontology. That baseline is the intended comparison point for the 2027 edition of this research.

SECTION 17

17 AI PM role archetypes

The seven archetypes in detail: what each one owns, what it screens for, what it pays, and who moves into it most easily.

8 archetypes with distinct skill profiles

Each has a different natural entry path

Archetype predicts requirements better than seniority alone

Choosing an archetype early is the highest-leverage decision a candidate makes, because it determines which 20% of the skill table actually matters for them.

Agentic AI PM

Owns autonomous or semi-autonomous workflows. Screens for tool-use design, failure recovery, permissioning and the organisational question of what a system is allowed to do unsupervised. 30.4% of postings · n=744 · median experience 5 yrs · median posted pay $195K · modal level Mid-level

MOST-REQUESTED SKILLS IN THIS ARCHETYPE	% OF ARCHETYPE POSTINGS
Artificial Intelligence	99.7%
AI Agents / Agentic AI	97.4%
Roadmapping	75.4%
Communication	71.1%
Prioritization	65.9%
Collaboration	64.4%
Ambiguity / Autonomy	60.1%
Stakeholder Management	59.3%

FOR JOB SEEKERS

The newest and least-defined. Accessible to anyone who can demonstrate they have designed a workflow where a system acts on a user's behalf and can fail safely.

GenAI / LLM Product PM

Owns user-facing products built on foundation models. Screens for prompt and context design, evaluation discipline, latency and cost trade-offs, and comfort with non-deterministic output. 25.9% of postings · n=635 · median experience 5 yrs · median posted pay $185K · modal level Mid-level

MOST-REQUESTED SKILLS IN THIS ARCHETYPE	% OF ARCHETYPE POSTINGS
Artificial Intelligence	99.4%
Large Language Models	72.8%
Roadmapping	72.8%
Communication	69.6%
Collaboration	66.9%
Stakeholder Management	64.6%
Prioritization	63.3%
Machine Learning	59.4%

FOR JOB SEEKERS

Most accessible to existing consumer or B2B SaaS PMs. The transferable asset is product judgement; the gap to close is evaluation and failure-mode literacy.

Data / ML Product PM

The oldest archetype. Owns ranking, recommendation, forecasting or personalisation. Screens for metrics rigour, experimentation and data-pipeline literacy. 13.3% of postings · n=326 · median experience 5 yrs · median posted pay $195K · modal level Mid-level

MOST-REQUESTED SKILLS IN THIS ARCHETYPE	% OF ARCHETYPE POSTINGS
Artificial Intelligence	98.2%
Roadmapping	81.3%
Communication	79.8%
Prioritization	77.6%
Collaboration	74.5%
Stakeholder Management	71.2%
Machine Learning	70.2%
Cross-functional Leadership	66.6%

FOR JOB SEEKERS

The natural destination for analytics, data and classical-ML product people. Metrics rigour is the currency.

AI Platform / Infra PM

Owns the internal platform other teams build on: model serving, inference cost, tooling, guardrails. Screens for systems thinking and developer empathy more than consumer product instinct. 8.5% of postings · n=208 · median experience 5 yrs · median posted pay $204K · modal level Mid-level

MOST-REQUESTED SKILLS IN THIS ARCHETYPE	% OF ARCHETYPE POSTINGS
Artificial Intelligence	100.0%
Roadmapping	84.1%
Prioritization	75.5%
Communication	70.2%
Collaboration	69.2%
Leadership	59.1%
Machine Learning	55.8%
Stakeholder Management	54.8%

FOR JOB SEEKERS

Most accessible to technical PMs, developer-tools PMs and engineers moving into product. Requires genuine systems fluency.

General AI Product Manager

Broad AI ownership without a named specialism, most common at companies adding their first AI PM. Screens for range. 8.0% of postings · n=197 · median experience 5 yrs · median posted pay $194K · modal level Mid-level

MOST-REQUESTED SKILLS IN THIS ARCHETYPE	% OF ARCHETYPE POSTINGS
Artificial Intelligence	95.9%
Roadmapping	59.4%
Collaboration	57.4%
Communication	56.9%
Stakeholder Management	54.8%
Prioritization	53.8%
Cross-functional Leadership	48.2%
Leadership	43.7%

FOR JOB SEEKERS

Typical of a company hiring its first AI PM. Rewards breadth and comfort with an unformed brief.

Applied AI / AI Applications PM

Embeds AI into an existing product surface. Screens for prioritisation and change management as much as AI depth. 7.6% of postings · n=187 · median experience 5 yrs · median posted pay $175K · modal level Mid-level

MOST-REQUESTED SKILLS IN THIS ARCHETYPE	% OF ARCHETYPE POSTINGS
Artificial Intelligence	100.0%
Roadmapping	81.3%
Communication	78.1%
Prioritization	77.5%
Collaboration	74.3%
Stakeholder Management	73.8%
Ambiguity / Autonomy	60.4%
Cross-functional Leadership	59.4%

FOR JOB SEEKERS

The most common internal transition: an existing PM whose surface gains an AI capability.

Conversational AI PM

Owns assistants, copilots and support automation. Screens for dialogue design, containment and deflection metrics, and integration with human escalation. 3.5% of postings · n=85 · median experience 5 yrs · median posted pay $187K · modal level Mid-level

MOST-REQUESTED SKILLS IN THIS ARCHETYPE	% OF ARCHETYPE POSTINGS
Artificial Intelligence	98.8%
Roadmapping	81.2%
Stakeholder Management	69.4%
Communication	61.2%
Prioritization	55.3%
Collaboration	55.3%
B2B / Enterprise Product	50.6%
Product Strategy	50.6%

FOR JOB SEEKERS

Accessible from support operations, CX and service design as well as product. Containment and escalation metrics are the language.

Responsible AI / Governance PM

Owns policy, safety, evaluation standards and regulatory readiness. Screens for governance frameworks, risk language and cross-functional authority. 2.7% of postings · n=66 · median experience 5 yrs · median posted pay $178K · modal level Mid-level

MOST-REQUESTED SKILLS IN THIS ARCHETYPE	% OF ARCHETYPE POSTINGS
Artificial Intelligence	100.0%
Roadmapping	93.9%
Responsible AI / AI Ethics	87.9%
Stakeholder Management	83.3%
Prioritization	75.8%
Communication	71.2%
Cross-functional Leadership	69.7%
Leadership	65.2%

FOR JOB SEEKERS

Accessible from risk, compliance, policy and legal backgrounds with product exposure. Concentrated in regulated sectors.

SECTION 18

18 What employers really want

The synthesis: reading the demand signal underneath the frequency tables.

Five conclusions the data supports directly

Three widely-held beliefs the data contradicts

Written for decisions, not for citation

Frequency tables describe what employers write. This section is about what they are buying.

Five conclusions the data supports

01 **They are buying product managers, not AI specialists.**
Across every cut of this dataset — seniority, geography, industry, archetype — product-craft skills are requested at rates at or above AI-specific skills. Product Strategy (53%) and Stakeholder Management (64%) outrank almost every technique in the modern AI stack. The AI layer is a qualifier on a product role, not a replacement for one.

02 **They are buying judgement under uncertainty.**
The distinguishing vocabulary of these postings is about deciding when a probabilistic system is good enough: evaluation, quality bars, guardrails, human review, failure handling. This is the capability that has no analogue in traditional product management and the one that interviews probe hardest.

03 **They are buying delivery, not exploration.**
Ownership language dominates: own the roadmap, ship the capability, define the metric, take the surface end-to-end. Research-adjacent language is rare. Employers have moved past experimentation budgets into shipped-product accountability.

04 **They are buying translation.**
Cross-functional leadership and influence without authority appear at high rates

because the AI PM sits between applied scientists, engineers, legal, risk and commercial teams. The job is substantially about making four groups agree on what ‘working’ means.

05 **They are buying commercial accountability.**
Business-layer skills appear throughout the higher-paying segments. Cost, pricing, ROI and unit economics separate the top quartile from the bottom more reliably than any technique. Inference cost has made economics a product concern rather than a finance one.

Three beliefs the data contradicts

A **“You need to be able to code.”**
84% of postings name no programming language at all. Python appears in 9% and SQL in 11%, and where they appear the surrounding language is analytical rather than engineering. The real technical bar is API and data literacy.

B **“You need an advanced degree.”**
22% state a degree as required, 53% never mention one, and 11% explicitly allow equivalent experience. A PhD appears in 2.3% of postings.

C **“You need to master the newest framework.”**
Named vendor tools are among the least-requested items in the ontology. The frontier techniques that dominate online discussion — RAG at 19.0%, fine-tuning at 7.1% — are minority requirements. Concepts transfer; tools churn.

The candidates who struggle are rarely the ones who know too little about models. They are the ones who cannot describe a product decision they owned end to end.

84%

of AI Product Manager postings name no programming language at all

The technical bar is real but widely misunderstood. It is API, data and cost literacy — enough to interrogate your own product — not production engineering.

PART V

The Playbook

The part to act on: what to learn, in what order, what to build, and how to position whatever background you are starting from.

19 The job-ready stack

20 Roadmaps

21 Portfolio

22 Before 2027

23 Takeaways

24 Methodology

25 Limitations

26 Appendix

SECTION 19

19 The job-ready AI PM stack

A prioritised learning stack derived from demand frequency, pay association, portability and scarcity — not from opinion.

Four tiers: must know, should know, differentiators, advanced

Every item ranked by an explainable Skill Opportunity Score

Built for sequencing, not for completeness

This is the section to act on. Everything before it describes the market; this converts it into an order of operations.

The Skill Opportunity Score

Ranking by raw frequency rewards skills that are so universal they cannot differentiate a candidate. Ranking by pay association rewards skills that merely correlate with seniority. The score used here combines five explainable components:

1. **Demand** (28%) — how often the skill appears, log-scaled so that the gap between 5% and 15% counts for more than the gap between 70% and 80%.
2. **Differentiation** (22%) — highest in the middle of the frequency range. A skill requested by almost every employer cannot separate one candidate from another; a skill requested by almost none is not yet worth the investment. The peak sits around a third of postings.
3. **Pay association** (20%) — the observed difference in median posted salary between postings that name the skill and those that do not, floored at zero and reported as association only.
4. **Cross-industry portability** (13%) — the share of industries in which the skill appears, rewarding capabilities that survive a sector change.
5. **Gate strength** (17%) — how often the skill appears in a required rather than a preferred block.

Skill Opportunity Score: what to learn first

Composite of demand frequency, pay association, cross-industry portability, relative scarcity and required-versus-preferred status

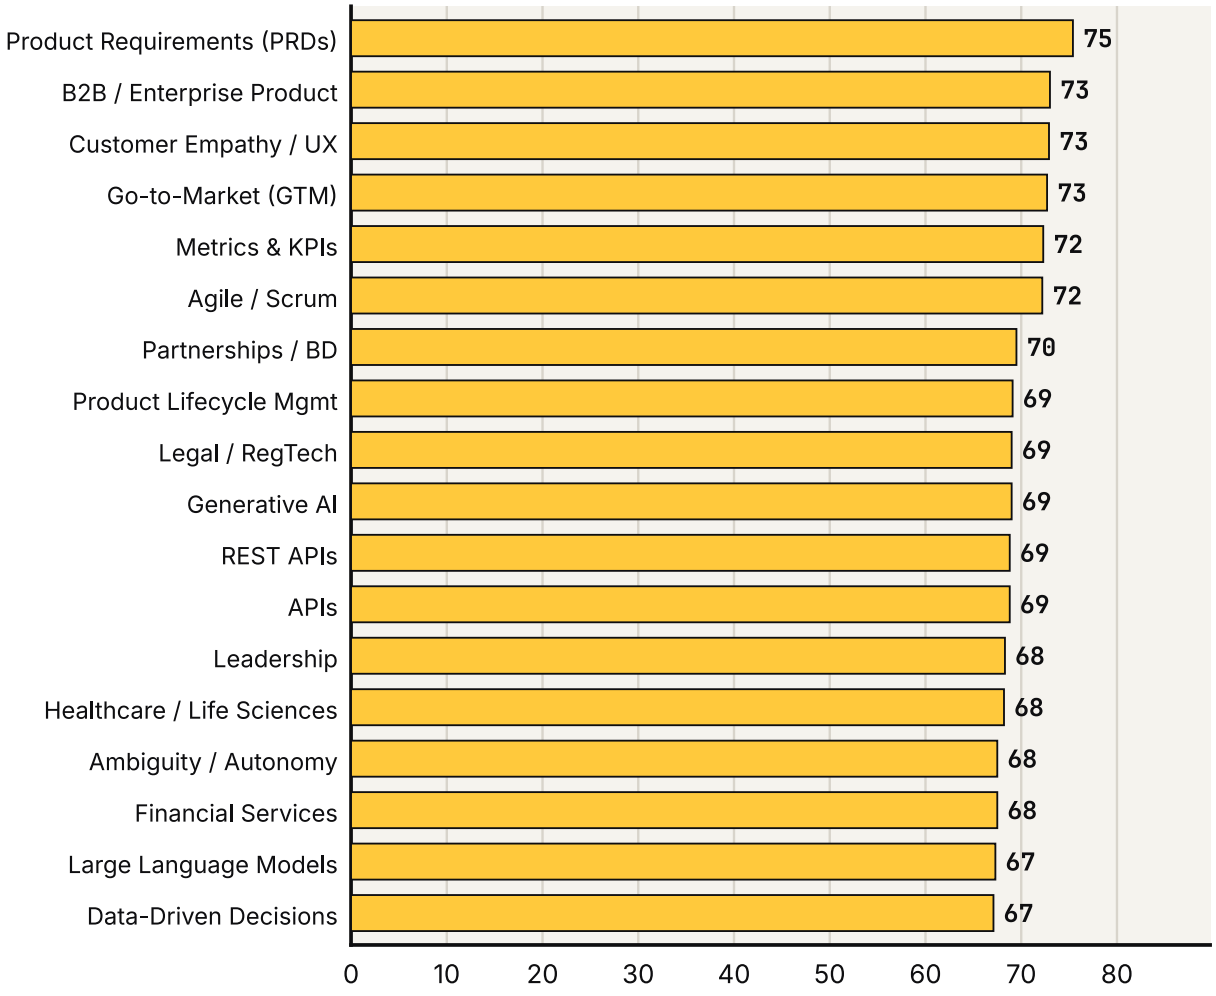

Product Requirements (PRDs) scores highest (75), appearing in 35% of postings — high-leverage skills are common enough to be expected and deep enough to stay scarce.

Source: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings (Tier A + Tier B), 2026 year-to-date, captured 12 September 2026. Score = 0.28 demand + 0.22 differentiation + 0.20 pay association + 0.13 portability + 0.17 gate strength, each normalised 0-1.

Figure 39 — Skill Opportunity Score: the composite ranking of what to learn.

#	SKILL	SCORE	% OF POSTINGS	DIFFERENTIATION	GATE STRENGTH	PAY SIGNAL
1	Product Requirements (PRDs)	75	34.8%	100	96	-7%
2	B2B / Enterprise Product	73	38.7%	92	70	+5%
3	Customer Empathy / UX	73	33.7%	97	84	-2%

#	SKILL	SCORE	% OF POSTINGS	DIFFERENTIATION	GATE STRENGTH	PAY SIGNAL
4	Go-to-Market (GTM)	73	26.5%	81	90	+5%
5	Metrics & KPIs	72	39.7%	90	88	-5%
6	Agile / Scrum	72	33.2%	96	82	-8%
7	Partnerships / BD	70	23.9%	75	75	+6%
8	Product Lifecycle Mgmt	69	25.6%	79	86	+2%
9	Legal / RegTech	69	34.0%	98	61	-2%
10	Generative AI	69	29.4%	88	77	-2%
11	REST APIs	69	28.3%	85	63	+4%
12	APIs	69	28.2%	85	80	-3%
13	Leadership	68	50.9%	65	81	+3%
14	Healthcare / Life Sciences	68	36.6%	96	56	-6%
15	Ambiguity / Autonomy	68	52.8%	60	80	+3%
16	Financial Services	68	39.9%	89	60	-5%
17	Large Language Models	67	46.2%	75	74	-2%
18	Data-Driven Decisions	67	24.4%	76	84	-3%
19	Mentorship / Coaching	67	21.2%	69	71	+6%
20	Adaptability / Learning	66	17.5%	61	76	+7%
21	Problem Solving	66	20.6%	68	89	-4%
22	Product Strategy	65	52.9%	60	78	+0%
23	Cross-functional Leadership	65	56.7%	52	85	-1%
24	Machine Learning	65	52.0%	62	73	-2%
25	AI Agents / Agentic AI	64	51.2%	64	67	+0%

Table 41 — Composite score, 0-100. Pay signal is the raw median difference and is reported for transparency, not as a causal estimate.

The four tiers

MUST KNOW — you will be screened out without these

These appear in a majority of postings and in required blocks. They are not differentiators; their absence is disqualifying.

Product Strategy (53%), Roadmapping (76%), Product Requirements (PRDs) (35%), Stakeholder Management (64%), Cross-functional Leadership (57%), Communication (71%), Metrics & KPIs (40%), Prioritization (67%), Artificial Intelligence (99%), Machine Learning (52%)

SHOULD KNOW — these decide the shortlist

Requested often enough to be expected in strong candidates, but not yet universal. This is where most candidates' preparation gap sits.

Generative AI (29%), Model Evaluation / Evals (20%), APIs (28%), SQL (11%), Experimentation & A/B Testing (20%), Product Discovery (10%), User Research (10%), Go-to-Market (GTM) (26%), Agile / Scrum (33%), Python (9%)

DIFFERENTIATORS — these move you up shortlists

Lower frequency, high signal. These are the skills that make a candidate memorable in a field where everyone claims AI fluency.

AI Agents / Agentic AI (51%), RAG (Retrieval-Augmented Generation) (19%), Guardrails & Safety (13%), Prompt Engineering (13%), Vector Search (5%), Embeddings (7%), Human-in-the-Loop (10%), Pricing & Monetization (12%), System Design / Architecture (11%), Platform / API Product (16%)

ADVANCED — depth for a specific archetype

Do not learn these first. Learn them once an archetype is chosen, because each is concentrated in a narrow band of roles.

Fine-tuning (7%), MLOps (6%), LLMOps (2%), Context Engineering (4%), Multimodal AI (3%), Feature Stores (1%), Model Context Protocol (MCP) (9%), LangGraph (2%), Synthetic Data (0%), Edge / On-device AI (6%)

Opportunity rankings by starting point

The same market looks different depending on where a candidate is standing. These rankings re-weight the score toward the segment of postings each group is realistically competing in.

Aspiring AI PMs (entry / mid)

#	SKILL	% OF SEGMENT POSTINGS	SEGMENT SCORE
1	Artificial Intelligence	99.1%	70
2	Stakeholder Management	64.3%	63
3	Roadmapping	67.8%	62
4	Communication	66.9%	62
5	Product Requirements (PRDs)	39.5%	62
6	Prioritization	62.1%	61
7	Collaboration	64.9%	61
8	Metrics & KPIs	39.5%	60
9	Ambiguity / Autonomy	46.9%	60
10	Agile / Scrum	38.6%	59
11	Large Language Models	46.5%	59
12	Cross-functional Leadership	49.4%	59

Existing PMs moving into AI

#	SKILL	% OF SEGMENT POSTINGS	SEGMENT SCORE
1	Artificial Intelligence	98.6%	70
2	Roadmapping	73.5%	64
3	Prioritization	69.6%	64
4	Leadership	55.3%	63
5	B2B / Enterprise Product	46.2%	63
6	Stakeholder Management	62.3%	62
7	Communication	68.2%	62
8	Cross-functional Leadership	56.5%	62
9	Ambiguity / Autonomy	51.5%	61
10	Collaboration	65.1%	61
11	Product Strategy	54.7%	61
12	Product Requirements (PRDs)	35.7%	60

Technical AI PMs

#	SKILL	% OF SEGMENT POSTINGS	SEGMENT SCORE
1	APIs	86.2%	75
2	Artificial Intelligence	99.2%	70
3	Roadmapping	75.2%	65
4	Communication	75.0%	65
5	Large Language Models	61.0%	65
6	Prioritization	70.5%	64
7	AI Agents / Agentic AI	64.5%	64
8	Stakeholder Management	65.0%	64
9	Ambiguity / Autonomy	56.8%	63
10	B2B / Enterprise Product	45.9%	63
11	Product Requirements (PRDs)	41.5%	62
12	Collaboration	67.1%	62

Senior / Principal PMs

#	SKILL	% OF SEGMENT POSTINGS	SEGMENT SCORE
1	Artificial Intelligence	99.1%	70
2	Roadmapping	81.2%	67
3	Prioritization	70.7%	65
4	Communication	72.8%	64
5	Ambiguity / Autonomy	58.0%	64
6	Cross-functional Leadership	61.6%	64

#	SKILL	% OF SEGMENT POSTINGS	SEGMENT SCORE
7	Leadership	54.9%	63
8	Stakeholder Management	61.6%	62
9	Collaboration	66.9%	62
10	Product Strategy	56.5%	62
11	B2B / Enterprise Product	43.4%	62
12	Machine Learning	54.9%	61

Executive AI Product leaders

#	SKILL	% OF SEGMENT POSTINGS	SEGMENT SCORE
1	Leadership	82.7%	74
2	Artificial Intelligence	99.6%	70
3	Roadmapping	87.6%	70
4	Product Strategy	75.2%	69
5	Communication	78.8%	66
6	Prioritization	74.8%	66
7	Stakeholder Management	70.8%	66
8	Cross-functional Leadership	65.0%	65
9	B2B / Enterprise Product	51.3%	65
10	Collaboration	73.5%	65
11	Financial Services	55.3%	63
12	Ambiguity / Autonomy	52.7%	62

SECTION 20

20 Roadmaps for aspiring AI Product Managers

Six starting points, six different sequences — each built from what the demand data says that background is missing.

Non-technical beginner, PM, engineer, data/ML, analyst, founder

Each roadmap targets the specific gap, not the whole curriculum

Sequenced against the experience distribution in Section 4

The fastest route into AI product management depends almost entirely on what you already have. These roadmaps are built backwards from the gap.

Non-technical beginner

You are missing the product core, not the AI. Given that only 8% of postings accept two years or fewer, treat this as a two-stage plan: get into product, then get into AI.

1. **Get product reps anywhere**

   Join a company as an associate PM, product analyst, support lead or operations analyst. The dataset is unambiguous that product experience is the binding constraint.

2. **Learn the product core properly**

   Discovery, PRDs, prioritisation, metrics. These are requested more often than any AI skill in this dataset.

3. **Build AI vocabulary, not AI mathematics**

   Understand what models do, how they fail, what they cost and how they are evaluated. You need to be credible in a room, not to train a model.

4. **Ship one AI feature end to end**

   Even internally. One shipped artefact outperforms any certificate — Section 21 lists what to build.

5. **Target the Applied AI archetype first**

   It is the most common internal transition and the least gated on prior AI experience.

Existing Product Manager

You have the hardest part already. Your gap is narrow, specific and closeable in a quarter of deliberate work.

1. **Close the evaluation gap first**

   Evaluations appear in 19.6% of postings and are almost never a nice-to-have. Learn to define a quality bar, build a golden set and decide what ships.

2. **Learn failure modes and guardrails**

   Hallucination handling, refusal behaviour, human-in-the-loop escalation. This is the vocabulary that signals you have shipped, not studied.

3. **Get comfortable with cost and latency**

   Inference economics now sit on roadmaps. Be able to discuss the trade-off between quality, cost and response time.

4. **Re-frame your existing work**

Most PMs have shipped something probabilistic — search, ranking, recommendations, fraud rules. Re-tell those stories in AI product language.

5 **Pick an archetype and go deep**

Section 17 shows which one your background maps to most cleanly.

Software Engineer

Your technical credibility is already above the bar. Your gap is demonstrating product judgement and commercial ownership.

1 **Stop leading with the stack**

The postings reward outcome ownership. Rewrite your experience around decisions and results, not systems built.

2 **Learn discovery and prioritisation formally**

These are the two skills hiring managers most often find missing in engineer-to-PM candidates.

3 **Own a metric, publicly**

Take responsibility for a number inside your current company before you apply for a role that requires it.

4 **Target platform and infrastructure archetypes**

AI Platform / Infra PM is the most natural landing point and rewards exactly the systems fluency you already have.

5 **Practise writing**

Executive communication appears throughout the senior and premium segments. Written clarity is the most under-invested engineering-to-PM skill.

Data or ML professional

You understand the models. Your gap is the customer and the commercial case.

1 **Learn the product craft explicitly**

Discovery, PRDs and roadmapping are requested far more often than modelling skills in these postings.

2 **Translate model metrics into product metrics**

The job is converting precision and recall into retention, cost and trust. Practise the translation.

3 **Talk to customers directly**

User research and customer interviews appear throughout the dataset and are the most common gap in data-to-product moves.

4 **Target Data / ML Product PM first**

It is the archetype where your existing credibility transfers with least discount.

5 **Build a commercial case**

Pricing, ROI and unit economics separate the higher-paying quartiles. Learn to argue for funding.

Business Analyst

You have stakeholder fluency and data literacy. Your gap is ownership and AI vocabulary.

1 **Move from requirements to outcomes**

Analysts document; PMs decide. Take a decision you can be wrong about.

2 **Deepen SQL and metrics design**

SQL appears in 11% of postings and is your strongest transferable asset.

3 **Learn how AI systems fail**

Evaluation, guardrails and human review. This converts analytical instinct into AI product judgement.

4 **Lead a cross-functional delivery**

Cross-functional leadership is requested in the large majority of postings.

5 **Target enterprise and regulated sectors**

Your process and governance fluency is a genuine advantage there.

Founder or entrepreneur

You have breadth, ownership and commercial instinct. Your gap is operating inside someone else's structure — and proving it.

1 **Translate founding work into product language**

Roadmaps, discovery, prioritisation, metrics. Name the artefacts explicitly; hiring managers screen for them.

2 **Show you can operate in a matrix**

Stakeholder management and influence without authority are requested constantly and are the most common founder objection.

3 **Pick a domain and commit**

Domain knowledge is a differentiator in this dataset and founders usually have one already.

4 **Target AI-native and early-stage employers**

The dataset shows these briefs reward range and independent judgement.

5 **Be specific about scale**

Senior postings ask for evidence of scope. Quantify users, revenue and team size honestly.

SECTION 21

21 Portfolio recommendations

Projects derived from what employers ask for — chosen because they produce evidence of the capabilities the demand data actually rewards.

Six projects, each mapped to specific requested skills

Each produces an artefact you can show, not a certificate

Ordered by return on effort for a candidate with limited time

A portfolio is not a demonstration that you can call an API. It is evidence that you can make a decision about a probabilistic system and defend it.

Build an evaluation harness for one real task

Pick a narrow task. Write 50-100 test cases with expected behaviour. Define what 'good' means before you measure. Run two models against it and write the ship/no-ship recommendation.

Evidences: Model Evaluation / Evals, Metrics & KPIs, Data-Driven Decisions, Product Requirements (PRDs)
Why it works: This is the single highest-return project in the dataset. Evaluations are requested widely and almost never as a nice-to-have, and almost no candidate arrives with one.

Write a real PRD for an AI feature

Take an existing product you know. Specify an AI capability for it: the user problem, the quality bar, the failure modes, the guardrails, the human-review path, the cost ceiling and the launch metric.

Evidences: Product Requirements (PRDs), Guardrails & Safety, Product Strategy, Human-in-the-Loop
Why it works: PRD-writing is requested across most of the sample and is the artefact hiring managers most often ask to see. An AI PRD that handles failure modes explicitly is rare.

Ship a retrieval-grounded assistant over your own documents

Index a real corpus. Make it answer questions with citations. Then do the hard part: measure groundedness, find where it fails, and write up what you changed.

Evidences: RAG (Retrieval-Augmented Generation), Embeddings, Vector Search, Hallucination Mitigation
Why it works: RAG appears in 19.0% of postings. Building one is common; measuring and improving one is not.

Design an agent with explicit permissions and failure recovery

Define a workflow where the system acts on a user's behalf. Specify what it may do unsupervised, what requires confirmation, how it recovers from a failed step, and how a human takes over.

Evidences: AI Agents / Agentic AI, Guardrails & Safety, System Design / Architecture, Human-in-the-Loop
Why it works: Agentic language is the fastest-moving category in the dataset, and the permissioning question is the one most candidates have never thought about.

Produce a build-versus-buy and unit-economics analysis

For one AI capability, model the cost per request at three volumes, compare an API vendor against a self-hosted option, and make a recommendation with a break-even point.

Evidences: Pricing & Monetization, P&L / Business Metrics, Business Strategy, Cloud Computing
Why it works: Commercial accountability separates the higher-paying quartiles more reliably than any technique, and almost no portfolio contains this.

Run a discovery study and publish what you learned

Interview eight to ten people about a workflow AI could change. Synthesise the findings. Write the opportunity assessment and the reason you would *not* build the obvious thing.

Evidences: User Research, Customer Interviews, Product Discovery, Prioritization
Why it works: Discovery skills are requested throughout the dataset and are the most common gap in technical candidates moving into product.

FOR JOB SEEKERS

Three finished projects with written decisions beat ten repositories. The artefact that gets a candidate hired is almost always the write-up, not the code: what you measured, what you found, what you decided, and what you would do differently.

SECTION 22

22 What to learn before 2027

Evidence-based priorities for the next twelve months, separated by how confident the data allows us to be.

High-confidence priorities grounded in current demand

Directional calls flagged explicitly as such

What to deprioritise, and why

Forecasting is where research of this kind usually stops being evidence. What follows is separated into what the data supports and what it merely suggests.

High confidence — grounded in observed demand

01 **Evaluation and quality measurement**
Requested in 19.6% of postings, almost always as a requirement rather than a preference, and present across every archetype and industry in the dataset. The most portable AI skill measured here.

02 **Product fundamentals, done properly**
Product Strategy at 53%, Roadmapping at 76%, Prioritization at 67%. These have not moved and will not. They are the base rate of the profession.

03 **Agentic system design**
The fastest-moving category in this dataset and the one where employer language is changing most quickly. Learning to specify what a system may do unsupervised is a transferable product skill, independent of framework.

04 **Cost and latency literacy**
Inference economics appear throughout the higher-paying segments. This is a permanent addition to the product manager's job, not a phase.

05 **Governance and responsible AI in regulated sectors**
Responsible AI appears in 23.2% overall but far higher in healthcare, financial services and public sector. Regulation is tightening, not loosening.

Directional — reasoned, not measured

READ THIS CAREFULLY

The three calls below are inferences from the direction of employer language in this dataset, not measured trends. They are labelled separately so readers can discount them appropriately.

A **Evaluation becomes a named discipline**
Evaluation language is already appearing as a distinct responsibility rather than a step inside delivery. Expect dedicated evaluation ownership in larger product organisations.

B **The agentic archetype separates from the GenAI archetype**
Agent postings already carry distinct vocabulary around permissioning, recovery and autonomy. That is how the Data/ML archetype separated from general product management a decade ago.

C **Context and memory design displaces prompt engineering as the named skill**
Prompt engineering appears as a component rather than a role. The broader framing — what the system knows, when, and at what cost — is where employer language is moving.

What to deprioritise

1. **Certifications.** Mentioned in 12.3% of postings, and no single certification family appears often enough to function as a signal.
2. **Deep model internals.** Transformer mathematics does not appear in these postings. Understanding behaviour and failure modes does.
3. **Framework tourism.** Named vendor tools sit at the bottom of the ontology by request frequency. Concepts transfer between them; the tools themselves churn annually.
4. **Breadth without evidence.** A CV listing twenty AI technologies with no shipped artefact performs worse than one project described in depth.

SECTION 23

23 Key takeaways

Ten sentences that carry the whole report.

01 AI product management is a product job with an AI qualifier — product craft is requested at or above the rate of every AI-specific skill in the dataset.

02 The median posting asks for 5 years of experience, and only 8% accept two years or fewer; this is not an entry-level field.

03 Degrees are rarely gates: 22% state one as required and 53% never mention education at all.

04 84% of postings name no programming language; the real technical bar is API, data and cost literacy.

05 Evaluation is the highest-return AI skill to learn — widely requested, rarely optional, and portable across every archetype and industry.

06 The frontier stack differentiates precisely because it is not yet standard: RAG at 19.0%, agents at 51.2%, fine-tuning at 7.1%.

07 Posted pay is high and wide: a global median of $193K against a 90th percentile of $265K.

08 What separates the highest-paying roles is scope — platform ownership, organisational leadership, commercial accountability — not tooling.

09 Demand is broad, not concentrated: 76% of hiring companies posted exactly one AI PM role, so the opportunity is wider than the obvious employer list.

10 Requirements are converging globally; what varies by market is seniority mix, working model and pay, not the substance of the job.

Learn the product craft. Add AI judgement. Ship something you can defend. In that order.

SECTION 24

24 Methodology

How the dataset was built, how postings were classified and de-duplicated, and how every measurement in this report was produced.

8,073 AI-relevant postings before de-duplication	2,875 duplicates removed	2,448 unique Tier A + Tier B postings in the core dataset

Every number in this report is reproducible from the dataset and scripts described here. Where a measurement is partial, its coverage rate is stated rather than implied.

1. Collection

Postings were collected directly from employer applicant-tracking systems wherever possible, and from public job-platform indexes where an ATS record was not reachable. Canonical employer records were preferred over aggregator copies throughout, because they carry requisition identifiers, structured compensation fields and publication dates that aggregators strip.

SOURCE	UNIQUE POSTINGS IN CORE DATASET	SHARE
LinkedIn (public job search)	1,816	74.2%
Greenhouse (employer boards)	244	10.0%
Ashby (employer boards)	185	7.6%
Lever (employer boards)	71	2.9%
Workday (employer career sites)	59	2.4%
The Muse	48	2.0%
SmartRecruiters (employer boards)	14	0.6%
Himalayas	9	0.4%
Jobicy	1	0.0%
Arbeitnow	1	0.0%

Table 47 — Source attribution after de-duplication; where the same requisition appeared in several sources, the canonical employer record was retained.

Applicant-tracking systems queried: Greenhouse, Lever, Ashby, SmartRecruiters, Workday, Rippling, Recruitee, Breezy and Workable public job-board endpoints. Job platforms and aggregators queried: LinkedIn public job search, The Muse, Arbeitnow, Himalayas, Jobicy, RemoteOK and Remotive. Employer board discovery was performed by probing candidate identifiers derived from public company registries and a large domain corpus, which is why the sample reaches well beyond the employers a keyword search would surface.

2. Scope and inclusion

The target population is product-management roles in which artificial intelligence, machine learning or generative AI is a core product responsibility. Two filters were applied in sequence.

1. **Role family.** A title classifier retains product-management roles (Product Manager, Product Owner, Head/Director/VP of Product, Group PM and equivalents) and excludes adjacent professions that share vocabulary: product marketing, product design, programme and project management, product analytics, sales engineering and engineering roles. These are different jobs and including them would distort every skill and salary figure in the report.

2. **AI relevance.** Each remaining posting is scored on AI centrality using signals drawn from the role-specific body of the description rather than the whole document.

Why the role body, not the whole description

Many employers now open job descriptions with a paragraph describing themselves as an AI-first company, and close with benefits and legal text. Scoring the full document causes ordinary roles at AI companies to be misclassified as AI roles. The classifier therefore isolates the role-specific body — from the first responsibilities or requirements heading to the start of the benefits or equal opportunity block — and measures AI centrality inside it.

Tiers

1. **Tier A** — AI, ML, GenAI, LLM or an equivalent term appears in the job title, with supporting AI content in the body.
2. **Tier B** — the title is a general product-management title, but AI is demonstrably a central responsibility: repeated AI ownership language in the role body, at sufficient density, in the requirements and responsibilities themselves.
3. **Tier C** — AI is present and meaningful but not central. Excluded from every headline figure in this report and reported separately.
4. **Tier D** — AI is incidental. Discarded.

The core dataset is Tier A + Tier B (2,448 postings: 1,784 Tier A and 664 Tier B). 2,750 Tier C postings were retained for sensitivity checking but excluded from all reported findings.

Tier A-only sensitivity check

Restricting to Tier A alone (1,784 postings) does not change the report's conclusions. Python moves to 9.9% (from 9.0%), SQL to 10.3% (from 10.7%), and degree-as-requirement to 22.3% (from 21.7%).

SKILL	TIER A ONLY	TIER A + B	DIFFERENCE
Artificial Intelligence	99.3%	99.1%	+0.2pp
Roadmapping	77.1%	76.1%	+1.0pp
Communication	71.8%	70.8%	+1.0pp
Collaboration	67.3%	66.7%	+0.6pp
Prioritization	66.6%	67.4%	-0.8pp
Stakeholder Management	64.1%	63.6%	+0.5pp
Cross-functional Leadership	56.8%	56.7%	+0.1pp
Machine Learning	54.1%	52.0%	+2.1pp
Ambiguity / Autonomy	53.3%	52.8%	+0.5pp
AI Agents / Agentic AI	52.4%	51.2%	+1.2pp
Product Strategy	52.2%	52.9%	-0.7pp
Leadership	49.3%	50.9%	-1.6pp
Large Language Models	46.2%	46.2%	+0.0pp
Financial Services	41.3%	39.9%	+1.4pp

Table 48 — Sensitivity of the top skill rankings to the Tier B inclusion decision.

3. De-duplication

A single requisition can appear on an employer's career site, two applicant-tracking mirrors and several aggregators. Counting those separately would inflate the sample and distort every percentage in this report. Six passes were applied in order, each linking a posting to an already-retained record rather than deleting it:

1. Identical canonical URL after query-string normalisation.
2. Identical full-description hash within the same normalised company name.
3. Identical full-description hash across different sources, catching syndicated copies where the company string differs.
4. Identical requisition identifier within the same normalised company.
5. Same company, normalised title, country and city, confirmed by description similarity above a shingle-overlap threshold.
6. Same company, normalised title and country, confirmed at a higher similarity threshold, collapsing multi-city variants of a single requisition.

Company names were normalised by removing legal suffixes and punctuation. Titles were normalised by stripping requisition numbers, location and working-model suffixes, gender-notation markers and abbreviations. Where two postings shared a company, title and location but carried distinct requisition identifiers and materially different descriptions, they were retained as separate requisitions.

DE-DUPLICATION STAGE	POSTINGS
AI-relevant postings identified (Tier A + B + C)	8,073
Duplicate records linked and removed	2,875
Unique postings retained (all tiers)	5,198
Core dataset — Tier A + Tier B unique postings	2,448

Table 49 — Counts at each stage of the de-duplication pipeline.

4. Field extraction

1. **Skills.** A 140-term ontology with regular-expression patterns for each term, covering synonyms, abbreviations, plurals and common spellings. Each skill is recorded once per posting regardless of repetition, and separately recorded as required or preferred where the description separates those blocks.
2. **Experience.** Candidate expressions are extracted with surrounding context, then scored. Expressions tied to the overall role ("years of product management experience", "years of relevant experience") are preferred over sub-requirements ("including two years in machine learning"), and references to company history or age are excluded.

3. **Education.** Degree level, field of study and requirement strength are read from the education sentence rather than the whole document, with explicit handling of "or equivalent experience" language.
4. **Salary.** Structured compensation fields from applicant-tracking systems are used first. Where absent, ranges are parsed from the description with currency, period and context validation, then annualised and converted to USD at spot rates on the capture date.
5. **Location and working model.** Structured location fields are used first, then parsed from the location string and description. Country names are canonicalised to a single spelling.
6. **Industry and archetype.** Assigned by weighted keyword classification of posting content.

5. Analysis

All percentages are shares of the core dataset unless stated otherwise. Where a field is only partially observed — salary and experience in particular — the denominator is the number of postings carrying that field, and the coverage rate is reported alongside. Salary comparisons across countries use nominal USD conversion and are not cost-of-living adjusted. Skill pay differences are computed within a single geography where sample size allows, in order to hold location broadly constant. A log-linear regression on 901 pay-disclosing postings models log posted-salary midpoint against skill indicators, country, seniority band, stated experience and working model, with heteroskedasticity-robust standard errors (adjusted R^2 0.48). Coefficients are reported as associations within this sample. Nothing in this report should be read as a causal estimate of the return on learning a skill.

6. Reproducibility

The following artefacts accompany this report: the cleaned posting dataset, the source ledger, the chart data, a written methodology note and the key findings summary. Every figure and table in this document is generated programmatically from the dataset, so the report can be regenerated end to end from the raw collection.

SECTION 25

25 Data limitations and coverage

What this dataset can and cannot support, stated plainly.

Every limitation below changes how a specific finding should be read. None of them are generic.

01 **This is a sample of observable postings, not a census.**
Public applicant-tracking endpoints and job-platform indexes do not cover every employer. Companies that recruit exclusively through agencies, internal mobility or closed networks are absent. The absolute counts in this report are a floor on market activity, not a total.

02 **Postings filled and de-listed before capture are invisible.**
The dataset observes roles that were live at the capture date. This under-counts earlier months of 2026 systematically, and is the reason the monthly velocity chart carries an explicit warning. It does not bias the skill, education or experience findings, which are measured within postings rather than across time.

03 **Salary is disclosed by 37% of postings and disclosure is not random.**
Pay-transparency legislation means United States, Canadian and some European postings disclose far more often than postings elsewhere. Every salary figure is therefore a statement about pay-disclosing postings, not about the market as a whole, and cross-country medians carry different coverage rates. The country table reports sample sizes so readers can weight accordingly.

04 **Posted salaries are not accepted salaries.**
These are advertised ranges, usually base pay, usually excluding equity and bonus. Self-reported compensation datasets measure a different and larger quantity. The two are not merged anywhere in this report.

05 **English-language postings are over-represented.**
Collection and classification operate on English text. Roles advertised solely in other languages are under-counted, which principally affects Japan, South Korea, Mainland China, Brazil and parts of continental Europe.

06 **Tier B classification is a judgement, implemented as a rule.**
Deciding when AI is ‘central’ to a general product role is inherently a threshold decision. The rule used here is documented in the methodology and its effect is quantified in the Tier A-only sensitivity table. Readers who prefer a stricter definition can read the Tier A column throughout.

07 **Skill measurement records mention, not depth.**
A posting naming Python is asking for something, but the ontology cannot distinguish ‘can read Python’ from ‘writes production Python’. Section 9 addresses

this qualitatively by reading the surrounding language; the quantitative figures should be read as mentions.

08 **Industry and archetype are inferred, not declared.**
Both are assigned by classifying posting content. Employers rarely state either explicitly. Misclassification is likely at the margins, particularly for diversified employers.

09 **Correlation is not causation, and the report never claims otherwise.**
Skill pay differences reflect the kind of role a skill appears in as much as the skill itself. The regression section exists specifically to show how much of the raw difference survives controls.

SECTION 26

26 Appendix

Full tables, taxonomies, the data dictionary and the source ledger.

A1 · Full skill ranking

#	SKILL	CATEGORY	POSTINGS	% OF SAMPLE
1	Artificial Intelligence	AI & ML Foundations	2,427	99.1%
2	Roadmapping	Product Management	1,863	76.1%
3	Communication	Leadership & Human Skills	1,734	70.8%
4	Prioritization	Product Management	1,650	67.4%
5	Collaboration	Leadership & Human Skills	1,632	66.7%
6	Stakeholder Management	Product Management	1,557	63.6%
7	Cross-functional Leadership	Product Management	1,389	56.7%
8	Product Strategy	Product Management	1,295	52.9%
9	Ambiguity / Autonomy	Leadership & Human Skills	1,292	52.8%
10	Machine Learning	AI & ML Foundations	1,273	52.0%
11	AI Agents / Agentic AI	Modern GenAI Stack	1,253	51.2%
12	Leadership	Leadership & Human Skills	1,246	50.9%
13	Large Language Models	AI & ML Foundations	1,130	46.2%
14	Financial Services	Domain Knowledge	976	39.9%
15	Metrics & KPIs	Product Management	972	39.7%
16	B2B / Enterprise Product	Product Management	947	38.7%
17	Healthcare / Life Sciences	Domain Knowledge	896	36.6%
18	Product Requirements (PRDs)	Product Management	852	34.8%

#	SKILL	CATEGORY	POSTINGS	% OF SAMPLE
19	Legal / RegTech	Domain Knowledge	832	34.0%
20	Customer Empathy / UX	Product Management	824	33.7%
21	Agile / Scrum	Product Management	813	33.2%
22	Generative AI	AI & ML Foundations	719	29.4%
23	REST APIs	Technical & Engineering	692	28.3%
24	APIs	Technical & Engineering	690	28.2%
25	Go-to-Market (GTM)	Product Management	648	26.5%
26	Product Lifecycle Mgmt	Product Management	627	25.6%
27	Data-Driven Decisions	Product Management	597	24.4%
28	Partnerships / BD	Business & Commercial	584	23.9%
29	Responsible AI / AI Ethics	Modern GenAI Stack	568	23.2%
30	Mentorship / Coaching	Leadership & Human Skills	518	21.2%
31	Problem Solving	Leadership & Human Skills	505	20.6%
32	Experimentation & A/B Testing	Product Management	492	20.1%
33	Model Evaluation / Evals	Modern GenAI Stack	479	19.6%
34	Executive Communication	Leadership & Human Skills	476	19.4%
35	Cybersecurity	Domain Knowledge	470	19.2%
36	RAG (Retrieval-Augmented Generation)	Modern GenAI Stack	466	19.0%
37	Adaptability / Learning	Leadership & Human Skills	429	17.5%
38	Sales Enablement	Business & Commercial	403	16.5%
39	Media / Entertainment	Domain Knowledge	399	16.3%
40	Platform / API Product	Product Management	399	16.3%
41	Public Sector / Defense	Domain Knowledge	392	16.0%
42	E-commerce / Retail	Domain Knowledge	387	15.8%
43	Business Strategy	Business & Commercial	374	15.3%
44	P&L / Business Metrics	Business & Commercial	368	15.0%
45	HR / Future of Work	Domain Knowledge	342	14.0%
46	Guardrails & Safety	Modern GenAI Stack	327	13.4%
47	Data Pipelines / ETL	Technical & Engineering	322	13.2%
48	Anthropic / Claude	Tools & Platforms	314	12.8%
49	Energy / Climate	Domain Knowledge	311	12.7%
50	Prompt Engineering	Modern GenAI Stack	309	12.6%
51	Dashboards & Reporting	Data & Analytics	297	12.1%
52	Storytelling	Leadership & Human Skills	294	12.0%
53	Data Analysis	Data & Analytics	287	11.7%
54	Pricing & Monetization	Business & Commercial	285	11.6%
55	System Design / Architecture	Technical & Engineering	266	10.9%
56	Market Research	Business & Commercial	265	10.8%
57	SQL	Technical & Engineering	262	10.7%

#	SKILL	CATEGORY	POSTINGS	% OF SAMPLE
58	Human-in-the-Loop	Modern GenAI Stack	256	10.5%
59	Cloud Computing	Technical & Engineering	249	10.2%
60	Product Discovery	Product Management	245	10.0%
61	Customer Obsession	Leadership & Human Skills	243	9.9%
62	Data Governance	Technical & Engineering	239	9.8%
63	AWS	Cloud Platforms	236	9.6%
64	User Research	Product Management	232	9.5%
65	Microsoft Azure	Cloud Platforms	228	9.3%
66	Competitive Analysis	Business & Commercial	227	9.3%
67	Manufacturing / Industrial	Domain Knowledge	227	9.3%
68	Influence w/o Authority	Leadership & Human Skills	221	9.0%
69	Python	Technical & Engineering	220	9.0%
70	Compliance & Regulation	Business & Commercial	214	8.7%
71	Model Context Protocol (MCP)	Tools & Platforms	213	8.7%
72	Budget / Resource Mgmt	Business & Commercial	211	8.6%
73	Supply Chain / Logistics	Domain Knowledge	198	8.1%
74	Jira / Confluence	Tools & Platforms	198	8.1%
75	OpenAI	Tools & Platforms	195	8.0%
76	Developer Tools	Domain Knowledge	195	8.0%
77	Embeddings	Modern GenAI Stack	173	7.1%
78	Fine-tuning	Modern GenAI Stack	173	7.1%
79	Critical Thinking	Leadership & Human Skills	163	6.7%
80	Customer Interviews	Product Management	163	6.7%
81	Big Data (Spark/Hadoop)	Technical & Engineering	155	6.3%
82	MLOps	Technical & Engineering	144	5.9%
83	CI/CD & DevOps	Technical & Engineering	143	5.8%
84	Edge / On-device AI	Technical & Engineering	137	5.6%
85	Automotive / Mobility	Domain Knowledge	136	5.6%
86	Hallucination Mitigation	Modern GenAI Stack	135	5.5%
87	Google Cloud (GCP)	Cloud Platforms	131	5.4%
88	Natural Language Processing	AI & ML Foundations	127	5.2%
89	B2C / Consumer Product	Product Management	115	4.7%
90	Vector Search	Modern GenAI Stack	113	4.6%

Table 50 — Complete normalised skill ranking, top 90 of 140 ontology terms.

A2 · Country sample sizes

COUNTRY	POSTINGS	SHARE	WITH SALARY	WITH EXPERIENCE
United States	1,283	52.40%	774	999
India	181	7.40%	0	136
United Kingdom	133	5.40%	13	63
Canada	104	4.20%	63	82
Germany	93	3.80%	0	18
France	59	2.40%	0	16
Singapore	58	2.40%	0	31
Israel	47	1.90%	0	37
Spain	42	1.70%	0	25
Netherlands	22	0.90%	0	13
Mainland China	20	0.80%	0	16
Australia	20	0.80%	0	0
Poland	20	0.80%	0	12
Portugal	19	0.80%	0	0
Ireland	19	0.80%	0	0
United Arab Emirates	16	0.70%	0	0
Brazil	15	0.60%	0	0
Malaysia	14	0.60%	0	0
Belgium	14	0.60%	0	0
Romania	13	0.50%	0	0
Vietnam	13	0.50%	0	0
South Korea	13	0.50%	0	0
Japan	13	0.50%	0	0
Hong Kong	12	0.50%	0	0
Philippines	11	0.40%	0	0
Mexico	9	0.40%	0	0
Taiwan	9	0.40%	0	0
Switzerland	9	0.40%	0	0
Greece	8	0.30%	0	0
New Zealand	8	0.30%	0	0

Table 51 — Sample sizes by country, including field coverage for the two partially-observed fields.

A3 · Title taxonomy

Search and classification covered the following title families and their common variants:

AI Product Manager · Product Manager, AI · Product Manager, Artificial Intelligence · AI/ML Product Manager · Machine Learning Product Manager · ML Product Manager · Generative AI Product Manager · GenAI Product Manager · LLM Product Manager · Agentic AI Product Manager · AI Agents Product Manager · Conversational AI Product Manager · AI Platform Product Manager · AI Infrastructure Product Manager · Applied AI Product Manager · Technical Product Manager (AI/ML) · Product Manager, Machine Learning · Product Manager, Generative AI · Senior AI Product Manager · Principal AI Product Manager · Lead AI Product Manager · Staff AI Product Manager · Group Product Manager (AI) · Director of AI Product · Director of Product (AI) · Head of AI Product · VP Product (AI) · AI Product Owner · Data Science Product Manager

General Product Manager and Technical Product Manager titles were additionally included wherever the role body made AI a central product responsibility (Tier B).

A4 · Skill ontology by category

AI & ML Foundations

Artificial Intelligence · Machine Learning · Large Language Models · Generative AI · Natural Language Processing · Recommendation Systems · Predictive Analytics · Foundation Models · Computer Vision · Deep Learning · Speech / Voice AI

Modern GenAI Stack

AI Agents / Agentic AI · Responsible AI / AI Ethics · Model Evaluation / Evals · RAG (Retrieval-Augmented Generation) · Guardrails & Safety · Prompt Engineering · Human-in-the-Loop · Embeddings · Fine-tuning · Hallucination Mitigation · Vector Search · Context Engineering · Multimodal AI · Synthetic Data

Technical & Engineering

REST APIs · APIs · Data Pipelines / ETL · System Design / Architecture · SQL · Cloud Computing · Data Governance · Python · Big Data (Spark/Hadoop) · MLOps · CI/CD & DevOps · Edge / On-device AI · Git / Version Control · Webhooks / Integrations · Model Inference / Serving · JavaScript / TypeScript · LLMOps · Data Modeling · Kubernetes · Microservices · Feature Stores · Docker · C++ / C# · Streaming (Kafka/Flink) · Java

Cloud Platforms

AWS · Microsoft Azure · Google Cloud (GCP)

Tools & Platforms

Anthropic / Claude · Model Context Protocol (MCP) · Jira / Confluence · OpenAI · Google Gemini · Tableau / Power BI / Looker · Snowflake · Excel / Spreadsheets · LangChain · Salesforce · Figma · Azure OpenAI · AWS Bedrock / SageMaker · LangGraph · Amplitude / Mixpanel · Notebooks (Jupyter) · PyTorch / TensorFlow · LlamaIndex · Hugging Face · MLflow / W&B · Weaviate / Qdrant / Milvus/ FAISS · Pinecone

Product Management

Roadmapping · Prioritization · Stakeholder Management · Cross-functional Leadership · Product Strategy · Metrics & KPIs · B2B / Enterprise Product · Product Requirements (PRDs) · Customer Empathy / UX · Agile / Scrum · Go-to-Market (GTM) · Product Lifecycle Mgmt · Data-Driven Decisions · Experimentation & A/B Testing · Platform / API Product · Product Discovery · User Research · Customer Interviews · B2C / Consumer Product · Product-Market Fit · Technical Documentation

Data & Analytics
Dashboards & Reporting · Data Analysis · Statistics · Cohort / Funnel Analysis

Business & Commercial
Partnerships / BD · Sales Enablement · Business Strategy · P&L / Business Metrics · Pricing & Monetization · Market Research · Competitive Analysis · Compliance & Regulation · Budget / Resource Mgmt

Leadership & Human Skills
Communication · Collaboration · Ambiguity / Autonomy · Leadership · Mentorship / Coaching · Problem Solving · Executive Communication · Adaptability / Learning · Storytelling · Customer Obsession · Influence w/o Authority · Critical Thinking

Domain Knowledge
Financial Services · Healthcare / Life Sciences · Legal / RegTech · Cybersecurity · Media / Entertainment · Public Sector / Defense · E-commerce / Retail · HR / Future of Work · Energy / Climate · Manufacturing / Industrial · Supply Chain / Logistics · Developer Tools · Automotive / Mobility · Education / EdTech · Telecom

A5 · Salary methodology

1. **Priority of sources.** Structured compensation fields published by the employer's applicant-tracking system are used first. Where absent, ranges are parsed from the posting text.
2. **Validation.** A parsed range is accepted only where a currency symbol or code is identifiable, the surrounding context is compensation-related, the pay period is determinable, and the annualised USD value falls within a plausible band.
3. **Annualisation.** Hourly figures are multiplied by 2,080, daily by 260, weekly by 52 and monthly by 12.
4. **Conversion.** Spot exchange rates retrieved on the capture date are applied. Rates are stored with the dataset so all conversions are reproducible.
5. **Midpoint.** Where a range is given, the midpoint is used. Where only one bound is given, that bound is used and flagged.
6. **Exclusions.** Equity, bonus, commission and benefits are excluded unless the posting's only stated figure was explicitly total compensation.

A6 · Data dictionary

FIELD	DEFINITION
job_id	Stable identifier assigned in this dataset
source / platform	Origin system for the retained canonical record
company / company_norm	Employer name as posted, and its normalised form used for de-duplication
raw_job_title / normalized_job_title	Title as posted, and its normalised form
role_family	Product Management (retained) or an excluded adjacent profession
ai_relevance_tier / ai_relevance_score	Tier A/B/C/D classification and its underlying score

FIELD	DEFINITION
seniority	Six-level band parsed from the title
role_archetype	One of the seven archetypes, assigned from title and responsibility language
industry	Employer sector, inferred from posting content
city / state / country / region	Normalised geography
work_mode	Remote, Hybrid, On-site or Not specified
min_experience_years / max_experience_years	Parsed experience requirement, with evidence string
degree_status / degree_levels / degree_fields	Education requirement strength, levels and fields
mba_requested / phd_requested / certifications	Credential flags
salary_min / salary_max / salary_currency / salary_period	Salary as posted
ann_salary_min / ann_salary_max	Annualised in the posting's own currency
usd_salary_min / usd_salary_max	USD-normalised at capture-date spot rates
salary_source	Structured ATS field or parsed from description
skills / skills_required / skills_preferred	Normalised ontology matches, split by requirement strength
posted_date / capture_date	First publication date where available, and the collection date
duplicate_cluster_id	Cluster identifier linking de-duplicated records

Table 52 — Fields present in the accompanying dataset export.

A7 · Source ledger and external references

Findings attributed to this research are computed from the posting dataset described above. External market-context sources consulted while framing the research questions, none of which contributed figures to any finding in this report, are listed below for readers who want independent cross-references on the wider AI labour market.

PUBLISHER	PUBLICATION	RELEVANCE
Stanford HAI	AI Index Report	Annual measurement of AI research, investment and labour-market indicators.
World Economic Forum	Future of Jobs Report	Employer-surveyed skill demand and job-growth projections.
LinkedIn Economic Graph	Jobs on the Rise; Skills on the Rise	Platform-derived role and skill growth rankings.
Lightcast	Public labour-market research	Job-posting analytics across occupations and skills.
Coursera	Job Skills Report	Learner-demand signals for technical and business skills.
Microsoft / LinkedIn	Work Trend Index	Survey research on AI adoption in the workplace.

PUBLISHER	PUBLICATION	RELEVANCE
Levels.fyi; Glassdoor; Built In; AmbitionBox	Self-reported compensation datasets	Cited in Section 10 as an alternative measurement of pay; not merged into this dataset.

Table 53 — External references consulted for context. No figure in this report is derived from them.

ABOUT THIS RESEARCH

The AI Product Manager Hiring Report 2026 is original research by Abhishek Ashtekar, published by ZeroToAIMastery.com. It is built on a purpose-collected dataset of 2,448 unique AI Product Manager job postings gathered from employer applicant-tracking systems and major job platforms, de-duplicated across sources and analysed with a 140-term normalised skill ontology. Data snapshot: 12 September 2026.

Citation: ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings, 2026.

AI PRODUCT MANAGER HIRING REPORT 2026

Original research by Abhishek Ashtekar, published by ZeroToAIMastery.com. Built on a purpose-collected dataset of 2,448 unique AI Product Manager job postings gathered from employer applicant-tracking systems and major job platforms, de-duplicated across sources and analysed with a 140-term normalised skill ontology.

HOW TO CITE

ZeroToAIMastery.com analysis of 2,448 unique AI Product Manager job postings, 2026.

ZeroToAIMastery.com

DATA SNAPSHOT / 12 September 2026