

Technical

Due Diligence

Merchantry Ltd

B2B trade supply marketplace · Series A · £6.2m round

CLIENT	Ashridge Ventures LLP
TARGET	Merchantry Ltd
ENGAGEMENT REF	TDD-2026-S1
SCOPE TIER	Full review — nine domains
FIELDWORK	4 – 15 May 2026
ISSUED	18 May 2026
VERSION	1.0 — Final
CLASSIFICATION	Sample — cleared for public distribution
PREPARED BY	Marc Roberts, Marcitect Solutions Limited

- OVERALL ASSESSMENT **ADEQUATE**
- RECOMMENDATION **PROCEED WITH CONDITIONS**

ABOUT THIS SAMPLE

Merchantry Ltd and Ashridge Ventures LLP do not exist. This report is a composite, written to demonstrate the structure, evidence standard and tone of a Marcitect due diligence deliverable. The findings reflect patterns I see repeatedly in Series A marketplaces; the figures are internally consistent but invented. No information from any real engagement appears in this document.

00 Basis, reliance and how to read this

0.1 Basis of assessment

Ten working days of fieldwork between 4 and 15 May 2026, with read-only access to source control, the cloud console and billing, the observability stack, and the delivery board. Six interviews totalling seven hours. Where an assessment rests on management assertion rather than verified evidence, it is marked LOW confidence and should be treated as a question to put to the company rather than a conclusion.

Materials reviewed

REF	DOCUMENT / SOURCE	RECEIVED	RELIED ON
M-01	Architecture overview and service map (Feb 2026)	2 May	Partly — see 2.1
M-02	Product roadmap to Q2 2027	2 May	Yes
M-03	Engineering org chart and hiring plan	2 May	Yes
M-04	AWS and third-party billing exports, Jun 2025 – Apr 2026	5 May	Yes
M-05	Management accounts, FY25 and YTD FY26	5 May	Yes
M-06	Employment and contractor agreements (14 documents)	7 May	Yes
M-07	Standard merchant supply agreement, v3 and v4	7 May	Yes
M-08	Penetration test report, Grosvenor Security, Oct 2025	8 May	Yes
M-09	Model evaluation notebook, catalogue matching (Nov 2024)	11 May	Yes

Access granted

- Source control: read-only across all 6 repositories, granted 4 May. Full history from first commit (March 2021), 18,400 commits.
- Cloud: read-only console and Cost Explorer across both production and staging accounts, granted 5 May.
- Observability and incidents: Grafana, and the incident channel archive from January 2025.
- Delivery: Linear workspace, read-only, full history.

Interviews

NAME	ROLE	DATE	DURATION
[Name]	CTO and co-founder	6 May	2 × 90 min
[Name]	Lead engineer, marketplace	8 May	75 min

NAME	ROLE	DATE	DURATION
[Name]	Engineer, pricing and credit	11 May	60 min
[Name]	Head of product	12 May	45 min
[Name]	Finance lead	13 May	45 min

Analysis performed

- Repository history analysis across 6 repositories and 18,400 commits — contribution distribution, change coupling, hotspot identification.
- Delivery metrics derived from version control and deployment records for the 12 months to 30 April 2026.
- Dependency and licence inventory across all repositories, with copyleft and end-of-life screening.
- Test coverage measurement, segmented by code path rather than reported globally.
- Cloud and third-party API cost decomposition by service and by cost driver, and unit cost-to-serve modelling at current, plan-year-one and plan-year-two volumes.
- Reconciliation of delivery costs against the cost-of-sales line in the management accounts.
- Architecture review against the growth assumptions stated in the Ashridge investment memo dated 21 April 2026.

Scope limitations

- No penetration testing was performed. Security findings rest on configuration review, the October 2025 third-party test, and interview.
- The merchant-facing Android application was excluded by agreement. It represents under 3% of order volume.
- Merchant supply agreements were reviewed for data rights only. Commercial terms were outside scope and sit with the client's legal advisers.
- Forecasting model performance could not be independently evaluated: the evaluation environment was not available within the fieldwork window. See DATA-01.

0.2 Reliance and limitations

This report has been prepared by Marcitect Solutions Limited for Ashridge Ventures LLP solely in connection with its evaluation of a proposed investment in Merchantry Ltd, and for no other purpose. It may be disclosed to Ashridge's professional advisers and to co-investors in the same transaction on the basis that they may read but not rely on it. No other party may rely on this report without our prior written consent.

Our work was performed between 4 and 15 May 2026 and reflects the position at that time. We have not updated it for subsequent events and are under no obligation to do so.

We have relied on information provided by Merchantry and its management without independent verification save where expressly stated. Where information was not provided, or was provided too late to examine, this is recorded under Scope limitations above.

Our total liability under or in connection with this report is limited as set out in our engagement letter dated 28 April 2026. Nothing in this report limits liability for fraud or for any other liability that cannot lawfully be limited.

This report is not

- an audit, or an opinion on financial statements;
- a penetration test, vulnerability assessment, or security certification of any kind;
- legal advice, including on intellectual property title, licence interpretation or data protection compliance — findings in those areas indicate where legal advice should be taken;
- a commercial, financial or product-market assessment;
- a guarantee that all defects, vulnerabilities or liabilities have been identified. Diligence is a sampling exercise conducted under time constraint.

0.3 Independence and conflicts

Marcitect Solutions Limited has no financial interest in Merchantry Ltd and no relationship with its management. Our fee is fixed and is not contingent on the outcome of the transaction or on the conclusions of this report. No conflicts have been identified. We do not accept diligence engagements in car sharing, car clubs, fleet or vehicle telematics.

0.4 How to read this report

Section 01 is the answer. Sections 02 and 03 are the evidence and the plan. If you read only two pages, read Section 01.

Domain assessments are relative to companies at a comparable stage and funding level, not to an enterprise standard. Finding severity is defined by consequence for this deal, not by technical seriousness in the abstract. Confidence markers distinguish what was observed directly from what rests on management assertion.

MARKER	MEANING
CRITICAL	Changes the decision, the price, or the structure. Resolve or protect against before close.
HIGH	Will consume material engineering capacity or capital within 12 months.
MEDIUM	Will constrain the team eventually. Fix within the investment period.
LOW	Noted for completeness. No action required for the investment case.

Remediation effort is stated in engineer-months of a mid-level engineer and converted at £11,000 fully loaded per engineer-month. Effort and elapsed time are stated separately. Assumptions are set out at 3.2.

01 The answer

1.1 Verdict

Merchantry's platform is capable of carrying the plan. The architecture is unremarkable in the way early-stage architecture should be, the team ships genuinely well, and nothing here threatens the investment case on its own. I would proceed.

But two things in the memo are wrong, and one of them is load-bearing. Gross margin is not 71%. Roughly £234,000 a year of delivery cost sits in an engineering budget line rather than in cost of sales; move it where it belongs and the restated margin is 58%. That matters more than the reclassification suggests, because the dominant cost driver scales with the number of merchants onboarded rather than with transactions — so margin deteriorates as the seller side of the marketplace grows, which is precisely what the plan proposes to do. Unremediated, I model gross margin at 54% by the end of the plan period against the 78% in the memo.

The fix is cheap and well understood: about two engineer-months of work to move the catalogue pipeline from nightly full reprocessing to incremental indexing. Total remediation across all findings is £198,000, or 3.2% of the round. Not doing the margin work costs roughly £680,000 a year by year two. This is the rare finding where the arithmetic is not close.

The second correction is the AI story. The demand forecasting product described in the deck is not in production and is not close; it runs as an internal weekly batch and reaches twelve pilot merchants as a PDF. Meanwhile the AI that is genuinely load-bearing — catalogue matching — barely appears in the materials, works well, and is not proprietary: it is a third-party embedding API applied to merchant-supplied catalogue data, under merchant agreements that do not clearly grant the rights being exercised. I would want that position confirmed in writing before close.

None of this is concealment. It is a company that has been moving quickly with a finance function that joined four months ago and a founder who describes the roadmap as though it has shipped. That is ordinary at this stage. It does mean the model needs rebuilding on restated numbers before final terms are agreed.

1.2 Thesis test

The assumptions below are taken from the Ashridge investment memo dated 21 April 2026. Each has been tested against the evidence available.

ASSUMPTION	VERDICT	BASIS	SEE
Platform supports 4× GMV within the investment period without re-architecture	SUPPORTED	Load headroom modelled from production telemetry; one database constraint noted	2.1, 2.3
Gross margin expands from 71% toward 78% as volume grows	NOT SUPPORTED	Margin misstated today; dominant cost scales with merchants, not transactions	2.3

ASSUMPTION	VERDICT	BASIS	SEE
Proprietary AI matching is a defensible advantage	PARTIALLY SUPPORTED	Matching works and is load-bearing, but is neither proprietary nor exclusive	2.6
AI demand forecasting is a near-term revenue line	NOT SUPPORTED	Not in production; no pricing, no contract, no delivery path in the roadmap	2.6, 2.8
Engineering can deliver the stated roadmap with four planned hires	PARTIALLY SUPPORTED	Roadmap requires ~26 engineer-months; realistic capacity is ~19	2.7, 2.8

1.3 Findings that matter

REF	SEVERITY	FINDING	EFFORT	COST	BITES
INF-01	CRITICAL	Gross margin misstated by 13 points; delivery costs sit outside COGS	—	—	Now
INF-02	HIGH	Catalogue pipeline cost scales with merchants, not transactions	2 em	£22k	At 300 merchants
DATA-01	HIGH	Forecasting product described as live is not in production	7 em	£77k	Now
DATA-02	HIGH	Merchant agreements do not clearly permit the model use being made of catalogue data	1 em	£11k	Now
TEAM-01	HIGH	Pricing and credit engine has a single author holding 78% of its commits	3 em	£33k	6–12 mo

1.4 What this means for the deal

Resolve or protect before close

- Rebuild the financial model on restated gross margin of 58%, agreed in writing with the company. The 78% exit margin in the memo is not achievable on the current cost structure and should not survive into the model. (INF-01)
- Obtain written confirmation, supported by legal advice, that the standard merchant supply agreement permits use of merchant catalogue data to train and improve matching models serving other merchants. If it does not, a v5 agreement and a re-consent exercise should be a condition. (DATA-02)
- Seek a specific warranty that customer-facing and investor-facing descriptions of the AI forecasting capability reflect what is in production, or agree the correction before the round is announced. (DATA-01)

Price into the round

- Nothing. The margin restatement corrects an input to the model rather than revealing a liability, and the remediation cost is immaterial against the round. If the model was built on 71%, however, the valuation derived from it needs revisiting on its own terms.

Fund in the first twelve months

- Incremental catalogue indexing, 2 engineer-months. The single highest-return piece of engineering work available to this company. (INF-02)
- Second engineer onto the pricing and credit engine, with documented handover. (TEAM-01)
- Matching accuracy monitoring, and a database read replica with a tested restore. (DATA-03, ARC-01)
- A decision on the forecasting product: fund it properly at 7 engineer-months, or remove it from the story. Do not leave it where it is. (DATA-01)

Accept and monitor

- Embedding provider concentration. Switching cost is about three engineer-weeks, which is tolerable. Revisit if provider pricing moves more than 25% or if inference exceeds 8% of net revenue. (VEND-01)
- The 2021 Polish agency assignment. Legal review recommended, but the code concerned is 4% of the current estate and largely superseded. (IP-01)

1.5 Remediation envelope

BUCKET	EFFORT	COST	ELAPSED
Before close	1 em	£11k	3 weeks
First 90 days	5 em	£55k	10 weeks
First 12 months	12 em	£132k	9 months
Total	18 em	£198k	—

Against a round of £6.2m, remediation represents 3.2% of new money and 14% of engineering capacity over twelve months, assuming the four planned hires land on schedule.

For comparison: the margin work alone costs £22,000 and is worth approximately £680,000 a year in retained gross profit by the end of the plan period.

02 Findings by domain

2.1 Product and architecture

ARC ASSESSMENT ADEQUATE

Does this design survive the next order of magnitude?

Assessment

A Rails monolith with two extracted services — catalogue processing and pricing/credit — on AWS, fronted by Postgres and Redis. This is the right architecture for a company of this size and I would have been more concerned to find microservices. Service boundaries follow genuine seams rather than fashion. Load testing in February 2026 sustained 4.1× current peak order volume with headroom, and the growth assumptions in the memo are comfortably within that. The architecture overview document is six months stale in two places, both noted to the CTO during fieldwork and neither material.

The one structural weakness is the database. A single Postgres instance carries the whole estate with no read replica, and reporting queries from the merchant dashboard contend with transactional load — visible in the p99 latency spikes every weekday morning. It is not urgent at current volume. It becomes urgent somewhere around 2.5× GMV, which the plan reaches in month fourteen.

REF	ARC-01
FINDING	The production database is a single instance with no read replica, and the restore procedure has never been tested.
EVIDENCE	Single db.r6g.2xlarge instance, no replica configured. Automated snapshots are running and retained for 7 days. No restore has been performed since the environment was created in 2022; the CTO confirmed this in interview. Reporting queries from the merchant dashboard execute against the primary — p99 latency rises from 180ms to 1.4s between 07:00 and 09:30 on weekdays.
IMPACT	Two exposures. Recovery time in a database failure is unknown and untested, which is a business continuity gap rather than a theoretical one. Separately, dashboard contention will force the replica work anyway at roughly 2.5× current GMV, reached in month 14 of the plan. Doing it under duress costs more than doing it deliberately.
REMEDIATION	Provision a read replica and route reporting and analytics queries to it. Perform a full restore into a scratch environment, time it, and document the RTO. Repeat the restore test quarterly.
EFFORT / COST	2 engineer-months · £22,000 · 4 weeks elapsed
MARKERS	SEVERITY MEDIUM MATERIAL TO THESIS N CONFIDENCE HIGH

What is working

Service boundaries are sensible and the team can explain why each one exists. Infrastructure is defined in Terraform and the definitions match what is deployed, which is less common at this stage than it should be.

2.2 Codebase and technical debt

CODE ASSESSMENT ADEQUATE

What does the next feature cost, and is that number rising?

Assessment

Coverage is 34% overall, which sounds poor and is not the useful number. Segmented by path, the order, payment and credit flows sit at 71%, and the untested 66% is concentrated in admin tooling, internal reporting and the merchant onboarding wizard. That is a deliberate allocation rather than neglect, and it is the allocation I would have chosen. A new engineer joining in February ran the system locally and shipped to production on day four, which is the best single proxy for codebase health available and it is a good result.

Hotspot analysis identifies the merchant onboarding wizard as the one area where high churn meets high complexity and low coverage. It has been rewritten twice and is the source of 31% of bugs logged in the last six months against 9% of commits. It is not a deal issue. It is the first thing I would fix if I were the CTO.

REF	CODE-01
FINDING	The merchant onboarding wizard concentrates a disproportionate share of defects and has no meaningful test coverage.
EVIDENCE	31% of bugs logged in Linear between November 2025 and April 2026 relate to onboarding, against 9% of commits touching that code. Coverage on the module is 11%. Two full rewrites in the repository history, most recently September 2025.
IMPACT	Onboarding is the growth-critical path: the plan more than triples merchant count. Defect load here consumes support time and directly slows the seller-side acquisition the plan depends on.
REMEDIATION	Characterisation tests around current behaviour before any further change. Then decompose the wizard into discrete steps with independent validation rather than a third rewrite.
EFFORT / COST	2 engineer-months · £22,000 · 6 weeks elapsed
MARKERS	SEVERITY LOW MATERIAL TO THESIS N CONFIDENCE HIGH

What is working

Dependencies are current — nothing more than one major version behind, no end-of-life runtimes within 18 months. Licence inventory is clean: 412 dependencies, all permissive, no copyleft anywhere in the production estate.

2.3 Infrastructure, reliability and cost to serve

INF ASSESSMENT WEAK

Does the platform stay up, and does it stay profitable as it grows?

Assessment

Reliability is fine. Uptime over the last twelve months is 99.94%, incidents are recorded properly, and the team finds out about problems before merchants do. Deployment is automated and rollback is tested rather than theoretical.

The economics are the problem, and they are the most significant finding in this report. Total delivery cost — AWS plus third-party inference — runs at £47,000 a month, or £564,000 annualised. Of that, £234,000 is booked to an engineering budget line rather than to cost of sales, because the infrastructure is managed by engineers and the classification was never revisited when the finance lead joined in January. Reported gross margin of 71% restates to 58% once those costs sit where they belong.

Underneath the classification issue is a structural one. The catalogue normalisation pipeline re-embeds the full catalogue nightly rather than processing changes incrementally — 1.4 million SKUs across 180 merchants, every night, regardless of what changed. Roughly 4% of SKUs change on a given day. The cost of this therefore scales with catalogue size, which is a function of merchant count, and has no relationship at all to transaction volume or revenue. Every merchant onboarded raises delivery cost from the day they sign, whether or not they ever transact.

Modelled forward on the plan's merchant acquisition curve, delivery cost grows 3.1× against net revenue growth of 2.7×, and gross margin falls to approximately 54% by the end of the plan period rather than expanding to 78%. Incremental indexing removes an estimated 62% of embedding spend and takes about two engineer-months.

REF	INF-01
FINDING	Gross margin is misstated by approximately 13 percentage points because delivery costs are classified outside cost of sales.
EVIDENCE	AWS and third-party API billing exports for the 12 months to April 2026 total £564,000. Reconciliation against the management accounts shows £330,000 in cost of sales and £234,000 in an engineering budget line covering compute, storage, data transfer and inference for the catalogue pipeline. These costs are consumed in delivering the service to customers and scale with usage rather than with headcount. Restating gives gross margin of 58.1% against the 71% reported.
IMPACT	The investment memo models exit margin from a 71% base. The base is wrong, and the direction of travel is wrong with it — see INF-02. Every downstream calculation in the model inherits the error.
REMEDIATION	Reclassify with the finance lead and restate FY25 and FY26 year to date. Adopt a standing definition of what enters cost of sales. This is an accounting correction, not engineering work.
EFFORT / COST	No engineering effort • finance time only • 2 weeks elapsed
MARKERS	SEVERITY CRITICAL MATERIAL TO THESIS Y CONFIDENCE HIGH

REF	INF-02
FINDING	The dominant delivery cost scales with merchant count rather than with transaction volume, so gross margin deteriorates as the marketplace grows its supply side.

EVIDENCE	The catalogue normalisation job re-embeds all 1.4m SKUs nightly. Change-rate analysis on the catalogue tables shows approximately 4% of SKUs modified per day. Embedding spend is £19,400 per month and has grown 9.4% month on month over the last six months against net revenue growth of 5.8%. Cost per active buyer account is £16.20 per month against net revenue per account of £52.
IMPACT	On the plan's merchant curve — 180 to approximately 600 merchants — delivery cost grows 3.1× against net revenue growth of 2.7×. Modelled gross margin at end of plan period is approximately 54% unremediated, against approximately 68% with incremental indexing. The difference is roughly £680,000 of annual gross profit in plan year two.
REMEDIATION	Move the pipeline to change-data-capture driven incremental embedding, processing only modified SKUs, with a weekly reconciliation pass. Cache embeddings by content hash so identical products across merchants are embedded once rather than 180 times.
EFFORT / COST	2 engineer-months · £22,000 · 5 weeks elapsed
MARKERS	SEVERITY HIGH MATERIAL TO THESIS Y CONFIDENCE HIGH

What is working

Operationally this is a well-run platform. 99.94% uptime, median incident recovery of 3.4 hours, no unplanned downtime exceeding 40 minutes in twelve months, and an incident log that is actually maintained. The cost problem is an architecture decision that was correct at 20 merchants and was never revisited — not a symptom of an undisciplined team.

2.4 Security and data protection

SEC ASSESSMENT **ADEQUATE**

What is the realistic exposure, and would anyone notice a breach?

Assessment

Better than typical for the stage. Secrets are in AWS Secrets Manager with no credentials anywhere in the repositories — I checked the full history, not just the current tree. The October 2025 penetration test found four medium and no high or critical issues, and all four were closed within seven weeks with evidence. Data is encrypted in transit and at rest. There is a data map, a retention policy that is actually applied, and processor agreements in place with the material subprocessors.

The weakness is joiners and leavers. There is no offboarding checklist, and access removal happens when someone remembers.

REF	SEC-01
FINDING	A contractor who left in August 2025 retained active write access to two repositories, and no offboarding process exists.
EVIDENCE	GitHub organisation audit shows one account with write access to the marketplace and catalogue repositories, last authenticated 14 March 2026, belonging to a contractor whose engagement ended 29 August 2025. No offboarding checklist exists in any

	documented form. Access removal for three previous leavers was performed between 3 and 11 weeks after departure, established from audit log timestamps.
IMPACT	Unauthorised code access and a live path to production for a party with no current relationship to the company. Also a control gap that will be raised by every enterprise customer security review and by SOC 2 if pursued, which the roadmap proposes for Q4.
REMIEDIATION	Revoke immediately. Audit all third-party and former-staff access across GitHub, AWS, the payment provider and the observability stack. Write a two-page joiner-mover-leaver checklist and tie it to the HR process.
EFFORT / COST	0.5 engineer-months · £6,000 · 2 weeks elapsed
MARKERS	SEVERITY HIGH MATERIAL TO THESIS N CONFIDENCE HIGH
REF	SEC-02
FINDING	Six of eleven engineers hold standing write access to the production environment.
EVIDENCE	AWS IAM review, 8 May. Six engineers hold roles permitting write operations against production data stores. No just-in-time elevation, no approval step, no session recording. Access was granted incrementally as the team grew.
IMPACT	Blast radius of a compromised laptop or credential is larger than it needs to be. Not currently exploited and no evidence of misuse; this is a control maturity gap rather than an active exposure.
REMIEDIATION	Move to time-bound elevation with an approval step for production write operations. Retain break-glass access for two named individuals.
EFFORT / COST	1 engineer-month · £11,000 · 3 weeks elapsed
MARKERS	SEVERITY MEDIUM MATERIAL TO THESIS N CONFIDENCE HIGH

What is working

Secrets management, encryption posture, remediation discipline after the penetration test, and UK GDPR documentation are all ahead of stage norms. The company has taken this seriously without over-investing in it.

2.5 IP, licensing and open source

IP ASSESSMENT **ADEQUATE**

Does the company own what it is selling?

Assessment

Chain of title is substantially clean. Both founders executed present-tense assignments in March 2021 before the first commit. All eleven current employees have assignment clauses in signed contracts. Six of seven contractors engaged since 2021 have executed assignments in the correct form. Neither

founder wrote core code while employed elsewhere; I checked employment dates against the repository history and the timeline is clear.

Open source exposure is negligible. 412 production dependencies, all permissive licences, no GPL or AGPL anywhere in the tree. This is unusually tidy and reflects a deliberate policy rather than luck.

One gap, and it is not serious. AI-assisted development is in use without a policy — approximately 35% of commits since mid-2025 carry assistance markers. Provider terms currently assign output to the user, so there is no ownership issue today, but the absence of a policy is worth closing.

REF	IP-01
FINDING	The 2021 development agency assignment is governed by Polish law and its scope is narrower than the current standard form.
EVIDENCE	Agreement with a Kraków-based agency, June to November 2021, governed by Polish law. The assignment covers "deliverables specified in the statement of work" rather than all work product, and does not address moral rights, which are non-waivable under Polish law. Repository analysis attributes approximately 4% of the current production estate to this engagement, largely in payment integration code that has since been substantially rewritten.
IMPACT	Low. The code concerned is a small and diminishing share of the estate and the counterparty relationship ended cleanly. It will nonetheless be raised in any future acquisition diligence, and is cheaper to resolve now than under transaction pressure.
REMEDIATION	Refer to Polish counsel for an opinion on the scope of the existing assignment. If confirmation is available cheaply, obtain a confirmatory deed. Otherwise document the position and monitor as the code is superseded.
EFFORT / COST	Legal cost only, estimated £2,000-£4,000 • 4 weeks elapsed
MARKERS	SEVERITY MEDIUM MATERIAL TO THESIS N CONFIDENCE MEDIUM

What is working

The licence position is genuinely clean, which removes the single most common source of late-stage surprise in software diligence. Founder assignments were done properly and done first.

2.6 Data and AI

DATA ASSESSMENT WEAK

Is the AI claim running in production, and does the company have the right to the data underneath it?

Assessment

There are two AI systems here. One is described extensively in the deck and is not in production. The other is barely mentioned, is genuinely load-bearing, works well, and is where the company's actual technical differentiation and its actual risk both sit.

Catalogue matching is the real one. Merchant product data arrives in inconsistent formats and the pipeline normalises it into a unified catalogue using embeddings from a third-party API, with a similarity threshold and a human review queue for ambiguous matches. It works — sampled against the review queue, precision holds around 94% on the cases I could evaluate. It is also not proprietary

in any meaningful sense: the embeddings are a commodity API call, and the differentiator is the accumulated review-queue corrections, which are a genuine asset but a modest one and not currently used to improve the model.

Demand forecasting is the one in the deck. It exists as a weekly batch job producing a spreadsheet, manually converted to PDF and emailed to twelve pilot merchants. There is no product surface, no pricing, no contract, and no line in the roadmap that would ship it. Described in the memo as "a near-term revenue line", it is approximately seven engineer-months from being a product.

REF	DATA-01
FINDING	The AI demand forecasting capability described as a near-term revenue line is not in production and has no delivery path in the roadmap.
EVIDENCE	The forecasting model runs as a weekly Python batch on a scheduled task, output written to S3 as CSV, manually converted and emailed to 12 pilot merchants. No API, no user interface, no entitlement or billing integration. The Q2 2027 roadmap contains no forecasting items. The model evaluation notebook is dated November 2024 and has not been re-run; the evaluation environment was unavailable during fieldwork, so model quality could not be independently assessed.
IMPACT	A revenue line in the plan that has no engineering path to existing. Either the plan overstates near-term revenue, or approximately 7 engineer-months of unbudgeted work is required — which is capacity the roadmap has already allocated elsewhere. See ROAD-01.
REMEDIATION	Decide. Either fund productionisation properly — API, entitlements, billing, and a re-evaluation of model quality before it is sold — or remove it from the plan and the customer-facing materials. The current position is the only one that is genuinely unhelpful.
EFFORT / COST	7 engineer-months • £77,000 • 5 months elapsed, if funded
MARKERS	SEVERITY HIGH MATERIAL TO THESIS Y CONFIDENCE HIGH
REF	DATA-02
FINDING	Merchant supply agreements do not clearly permit the use being made of merchant catalogue data to improve matching for other merchants.
EVIDENCE	Standard merchant supply agreement v3 (in force for 141 of 180 merchants) and v4 (39 merchants). Both grant Merchantry a licence to use catalogue data "for the purpose of listing and promoting the Merchant's products on the Platform". Neither addresses model training, model improvement, or use of one merchant's data in systems serving competitors. The review-queue corrections corpus, which management describes as a defensible data asset, is built from exactly this data.
IMPACT	The data asset described in the memo as a source of defensibility may not be one the company is contractually free to exploit. Merchants on this platform compete with each other, so the objection is a live commercial one rather than a theoretical legal one. Remediating later requires re-consent across the merchant base, which is materially harder at 600 merchants than at 180.

REMEDIATION	Obtain legal advice on the current position. Issue a v5 agreement with explicit and fairly-worded data rights, and run a re-consent exercise now while the merchant base is small. Until resolved, hold the corrections corpus separately and do not train on it.
EFFORT / COST	1 engineer-month · £11,000 plus legal · 8 weeks elapsed
MARKERS	SEVERITY HIGH MATERIAL TO THESIS Y CONFIDENCE MEDIUM

REF	DATA-03
FINDING	Matching model accuracy has not been measured since launch in November 2024.
EVIDENCE	Single evaluation notebook dated 14 November 2024 reporting 96.2% precision on a held-out set. No monitoring, no scheduled re-evaluation, no alerting on match confidence distribution. Human review queue volume has risen 34% over six months against catalogue growth of 22%, which is consistent with degrading match quality but has other possible explanations.
IMPACT	Match quality drives merchant satisfaction and buyer search relevance, both core to the marketplace. Degradation would currently be invisible until merchants complained. The review queue divergence is a weak signal that it may already be occurring.
REMEDIATION	Instrument match confidence distribution and review-queue rate as monitored metrics with alerting. Re-run the evaluation against a fresh held-out set. Schedule quarterly re-evaluation.
EFFORT / COST	1.5 engineer-months · £17,000 · 4 weeks elapsed
MARKERS	SEVERITY MEDIUM MATERIAL TO THESIS N CONFIDENCE HIGH

What is working

Catalogue matching is a well-chosen application of the technology to a real and expensive problem, and it works. The human review queue is the right design decision and shows the team understood the failure mode. Inference costs are high but the capability earns its place — the finding at INF-02 is about how it is run, not whether it should exist.

2.7 Engineering team and delivery capability

TEAM ASSESSMENT ADEQUATE

Can these people build the thing the plan requires?

Assessment

Eleven engineers and a CTO who is technically credible and has stopped being the best engineer in the room, which is the correct trajectory. Delivery is genuinely good: 4 to 6 deployments a week, median change lead time of 2.1 days, change failure rate around 9%, median recovery 3.4 hours. Those are respectable numbers for a team this size.

One caveat on the delivery metrics. Approximately 35% of commits since mid-2025 carry AI assistance markers, and throughput measures are less comparable across that boundary than they

appear. I have weighted the qualitative evidence — what shipped against what was promised six months ago — more heavily as a result, and it corroborates the picture: of eleven roadmap items committed for H2 2025, eight shipped, two slipped a quarter, one was dropped deliberately. That is a team that estimates honestly.

The concentration risk is specific and named.

REF	TEAM-01
FINDING	The pricing and credit engine has effectively one author, holding 78% of commits, and no other engineer has worked in it in over a year.
EVIDENCE	Contribution analysis across the pricing-credit repository: one engineer accounts for 78% of commits and 84% of lines changed since inception. The second-highest contributor last committed in March 2025. No design documentation. In interview, two other engineers described the credit limit logic as something they would not want to change without help.
IMPACT	The credit engine sets buyer credit limits and underwrites payment terms; it is the mechanism by which the marketplace takes balance sheet risk. An unplanned departure would leave the highest-consequence system in the business without an owner who understands it, at a point in the plan where credit exposure is growing.
REMEDIATION	Pair a second engineer into the module on live work, not as a documentation exercise. Write down the credit decisioning rules and their rationale. Target a second engineer able to make a substantive change unsupervised within one quarter.
EFFORT / COST	3 engineer-months • £33,000 • one quarter elapsed
MARKERS	SEVERITY HIGH MATERIAL TO THESIS Y CONFIDENCE HIGH

What is working

Delivery discipline is real and the team is honest about what it can do. Technical decisions are recorded — there are 23 decision records in the repository, current and actually referenced. Attrition is one engineer in three years.

2.8 Roadmap feasibility

ROAD ASSESSMENT **WEAK**

Does the plan fit the team?

Assessment

The roadmap to Q2 2027 contains work I estimate at 26 engineer-months. Available capacity over the same period, loading honestly for support, incidents, interviews and the debt the team already carries, is approximately 19 engineer-months — and that assumes all four planned hires land on schedule, which at current time-to-hire of 11 weeks means the last of them contributes for barely half the period.

The gap is 7 engineer-months, and it is almost exactly the forecasting work at DATA-01 that the roadmap does not contain. That is unlikely to be a coincidence: the roadmap reflects what the team intends to build, and the memo reflects what the company intends to sell, and the two documents have not been reconciled.

REF	ROAD-01
FINDING	The roadmap requires approximately 26 engineer-months against realistic capacity of 19, before any remediation work is added.
EVIDENCE	Roadmap items to Q2 2027 estimated independently and cross-checked against comparable completed work in the repository history. Capacity modelled from current headcount plus four planned hires at 11 weeks median time to hire, loaded at 70% for support, incidents and recruitment. Adding the 18 engineer-months of remediation in this report takes the requirement to 44 against 19.
IMPACT	Either the roadmap slips, the remediation does not happen, or the hiring plan expands. The plan does not currently acknowledge the choice. Given that the highest-value item available is the two-month margin fix at INF-02, the sequencing of this decision matters more than the size of the gap.
REMEDIATION	Reconcile the roadmap with the memo in a single planning session with the CTO and the board. Sequence remediation ahead of new feature work in the first quarter — INF-02 pays for a substantial part of the rest.
EFFORT / COST	Planning exercise · no engineering cost
MARKERS	SEVERITY HIGH MATERIAL TO THESIS Y CONFIDENCE MEDIUM

What is working

The roadmap itself is well constructed and the estimates within it are honest. The problem is the gap between it and the commercial narrative, not the quality of the planning.

2.9 Third parties and vendor dependency

VEND ASSESSMENT **ADEQUATE**

What happens when a supplier changes the terms?

Assessment

Three vendors are genuinely load-bearing: AWS, the payment and credit provider, and the embeddings provider. The payment relationship is contracted to March 2028 with a defined price schedule and no adverse renewal terms, and migration would be painful but is understood. AWS concentration is normal and unremarkable at this stage.

REF	VEND-01
FINDING	The embeddings provider is accessed directly with no abstraction layer, and inference spend is growing faster than revenue.
EVIDENCE	Provider SDK called directly from 14 call sites across two repositories. No interface abstraction, no evaluation of alternatives since selection in 2024. Month-to-month contract with no committed pricing. Inference spend is 12.9% of net revenue and growing 9.4% month on month against revenue growth of 5.8%.
IMPACT	Exposure to unilateral pricing changes on a cost line that is already a material share of revenue. Switching cost is modest — approximately 3 engineer-weeks to introduce an

	abstraction and re-evaluate — but it is 3 weeks the team does not have on a day when the provider changes terms.
REMEDIATION	Introduce a thin abstraction over the embedding calls and benchmark two alternatives against the current provider on the existing evaluation set. Do this alongside the INF-02 work, which touches the same code.
EFFORT / COST	0.75 engineer-months · £8,000 · alongside INF-02
MARKERS	SEVERITY MEDIUM MATERIAL TO THESIS N CONFIDENCE HIGH

What is working

The payment and credit provider relationship is well contracted and the terms were negotiated with attention. There is no vendor here that is also a competitor.

03 Remediation and plan

3.1 Remediation register

REF	ACTION	SEVERITY	EFFORT	COST	OWNER
BEFORE CLOSE					
INF-01	Restate gross margin and rebuild the model on 58%	CRITICAL	—	—	Finance lead / Ashridge
DATA-02	Legal opinion on merchant data rights; hold corrections corpus	HIGH	1 em	£11k	Counsel / CTO
DATA-01	Correct customer-facing forecasting claims or warrant them	HIGH	—	—	CEO / counsel
SEC-01	Revoke dormant contractor access	HIGH	—	—	CTO
FIRST 90 DAYS					
INF-02	Incremental catalogue indexing and content-hash caching	HIGH	2 em	£22k	CTO
SEC-01	Full third-party access audit and JML checklist	HIGH	0.5 em	£6k	CTO
TEAM-01	Second engineer into pricing and credit, on live work	HIGH	1.5 em	£17k	CTO
ROAD-01	Reconcile roadmap with commercial plan	HIGH	—	—	CTO / board
VEND-01	Abstraction over embedding provider, benchmark alternatives	MEDIUM	0.75 em	£8k	Catalogue team
DATA-02	v5 merchant agreement and re-consent exercise	HIGH	0.25 em	£2k	Counsel / commercial
FIRST 12 MONTHS					
DATA-01	Productionise forecasting, or formally drop it	HIGH	7 em	£77k	CTO / product
ARC-01	Read replica, query routing, tested restore	MEDIUM	2 em	£22k	Platform
TEAM-01	Complete handover; credit rules documented	HIGH	1.5 em	£16k	CTO
DATA-03	Match quality monitoring and re-evaluation	MEDIUM	1.5 em	£17k	Catalogue team

REF	ACTION	SEVERITY	EFFORT	COST	OWNER
SEC-02	Just-in-time production access	MEDIUM	1 em	£11k	Platform
ACCEPT AND MONITOR					
IP-01	Polish agency assignment — legal opinion, then monitor	MEDIUM	—	£2–4k	Counsel
CODE-01	Onboarding wizard — characterisation tests when next touched	LOW	—	—	CTO

3.2 Cost model and assumptions

- Effort is stated in engineer-months of a mid-level engineer at £11,000 fully loaded per month, inclusive of employer costs, tooling and management overhead. This is consistent with Merchantry's actual engineering cost per head.
- Effort and elapsed time are stated separately. Work dependent on a hire carries the 11-week median recruitment lead time in its elapsed figure.
- Estimates assume the current team plus the four hires already in the plan, and no specialist contractors.
- Estimates carry a range of roughly $\pm 40\%$. They are sized for budgeting, not as a fixed-price quotation.
- The £680,000 figure for retained gross profit at INF-02 is modelled on the plan's own merchant acquisition and revenue curve, with embedding cost held at the observed per-SKU rate and a 62% reduction assumed from incremental indexing. It is sensitive to the merchant curve: at 70% of planned merchant acquisition the figure is approximately £460,000.
- Legal costs are indicative only and are not within our competence to estimate precisely.

3.3 First hundred days

WHEN	ACTION	WHY IT IS FIRST
Weeks 1–2	Access audit and JML checklist. Restated margin agreed and circulated to the board.	Both are cheap, both are embarrassing to be caught by later, and neither needs a plan.
Weeks 3–8	Incremental catalogue indexing shipped. Embedding abstraction alongside it.	Highest return available. Pays for most of the rest of this list within two quarters and touches the same code as VEND-01.
Weeks 4–13	Second engineer into pricing and credit. Credit rules written down.	Takes a full quarter to be real, so it has to start early. It cannot be compressed later.
Weeks 6–10	Roadmap reconciled with commercial plan. Forecasting funded or dropped.	Needs the board, and needs to happen before the next planning cycle sets expectations again.

3.4 Questions to put to the company

- What did the merchants in the forecasting pilot understand they were receiving, and has any commercial commitment been made to any of them? Nothing in the materials suggests it has, but the pilot has run for seven months and the exposure would be commercial rather than technical.
- Was the decision to run nightly full catalogue reprocessing ever revisited, and if not, what is the mechanism by which cost decisions get reviewed? The answer tells you more about how this team will behave at 600 merchants than the finding itself does.
- What is the CTO's own view of the roadmap gap at ROAD-01? He estimated the roadmap honestly, so he may already know. Whether the gap has been raised with the board is the useful question.
- Has any merchant ever objected to catalogue data being used to improve matching? A single prior objection changes the risk profile at DATA-02 considerably.

3.5 Suggested deal protections

Offered as technical input to Ashridge's legal advisers. This is not legal advice and the drafting is a matter for counsel.

MECHANISM	COVERING	RATIONALE
Condition precedent	INF-01	Restated management accounts on the agreed cost-of-sales definition, delivered before completion. The model that supports the valuation should be built on numbers both sides have signed.
Warranty	DATA-01	That customer-facing and investor-facing descriptions of product capability reflect what is in production at completion. Cheap to give if true, and it forces the conversation now rather than at Series B.
Specific indemnity	DATA-02	Against claims arising from use of merchant catalogue data in model training prior to completion. Narrow, time-limited, and likely to fall away once the v5 agreement is in place.
Board reporting undertaking	INF-02, ROAD-01	Gross margin on the restated basis, and roadmap-versus-capacity, reported at every board meeting for the first year. These are the two things that will drift quietly.

04 Appendices

A Evidence log

Documents M-01 to M-09 as listed at 0.1, with the finding references each supports. Cloud billing exports, repository analysis outputs and coverage reports supplied separately as machine-readable files.

B Interview record

Five participants, six sessions, seven hours in total between 6 and 13 May 2026. Notes are held on file and available to Ashridge on request.

C Method and tooling

Repository analysis performed with git-of-theseus and custom scripts over full history; outputs supplied. Coverage measured with SimpleCov, segmented by path rather than reported globally. Dependency and licence inventory generated with Syft, cross-checked manually for the 40 most-depended-upon packages. Cloud cost decomposition from AWS Cost and Usage Reports at resource granularity, joined to service ownership.

Delivery metrics were derived from deployment records and version control rather than self-reported. Note the limitation recorded at 2.7: approximately 35% of commits since mid-2025 carry AI assistance markers, and throughput measures are not directly comparable across that boundary. Qualitative delivery evidence has been weighted accordingly.

D Scan outputs

Dependency inventory (412 production packages, all permissive), vulnerability summary, coverage report by path, and repository contribution analysis. Supplied separately as merchantry-scans-20260518.zip.

E Open items

REF	ITEM	RAISED	STATUS
O-01	Forecasting model evaluation environment — access not available within fieldwork window	11 May	Open
O-02	Merchant agreement v2 (pre-2023) — 4 merchants remain on this form, copies not located	12 May	Open
O-03	Contractor assignment, one 2023 engagement — signed copy not produced	13 May	Answered 15 May; executed copy provided

F Glossary

Cost to serve	The infrastructure and third-party cost consumed in delivering the product to a customer. Belongs in cost of sales.
Embedding	A numerical representation of text that allows similar items to be matched by proximity. Here, used to match merchant product listings to a unified catalogue.
Engineer-month	One mid-level engineer working for one month. The unit used throughout this report for remediation effort.
Incremental indexing	Processing only records that have changed, rather than reprocessing everything on a schedule.
Read replica	A copy of a database that serves read queries, so that reporting does not compete with live transactions.

Marcitect Solutions Limited · Technical due diligence and fractional CTO · marcitect.com

This is a composite sample. Merchantry Ltd and Ashridge Ventures LLP do not exist and no real engagement is described.