



GADIRLABS · SAMPLE DELIVERABLE

FICTIONAL COMPANY · ILLUSTRATIVE DATA

Asterpath production-readiness audit

A senior-engineer assessment of a fictional AI-assisted planning product: critical journeys, release evidence, operating costs and the conditions for a responsible launch.

COMPANY

Asterpath Ltd

PRODUCT

Asterpath web application

ASSESSMENT DATE

7 September 2026

DECISION

Do not launch yet

Asterpath Ltd, its product, repositories, evidence and findings are entirely fictional. Any resemblance to a real business is coincidental. This document demonstrates the format and reasoning of a gadirlabs audit. It is not client work, a testimonial or evidence that these events occurred.

EXECUTIVE DECISION

Do not launch yet

Fix and retest the paid-access recovery path before release.

The product has credible foundations and the main journeys work in staging. A reproduced payment-to-entitlement failure is a launch blocker. AI-cost reconciliation and recovery-path automation should follow before paid acquisition scales.

4**critical journeys**

Agreed before review.

3**material findings**

One High, two
Medium.

202**automated checks**

None verifies payment
recovery.

1**launch blocker**

Paid-access recovery.

What must happen before release

1. Make payment-event failure retryable and add durable reconciliation.
2. Prove the affected journey with a staging failure/recovery test.
3. Attach the passing checks and tested staging build to one immutable release record.

What can follow before growth

- Reconcile AI-provider invoices against product usage, including retries and failed attempts.
- Add automated unhappy and recovery coverage for the remaining revenue and AI journeys.
- Retest only the agreed findings and update the decision against a named release candidate.

DECISION BOUNDARY

"Do not launch yet" applies only to the fictional reviewed version, four journeys, staging environment and read-only consoles described here. It is not a guarantee of security, compliance, uptime or defect-free software.

DIMENSIONAL SCORECARD

Eight scores, each bounded by evidence

Scores make strengths and weak areas comparable without pretending the product has one meaningful “readiness percentage”. Each integer is an evidence-backed maturity level for the fictional reviewed scope; confidence describes the evidence, not the quality.

Architecture & ownership

Clear service ownership; failure continuity is only partly exercised.



3/5 Medium confidence

Critical journeys & UX

Four journeys run; alternate and recovery UX is incomplete.



3/5 High confidence

Security & data protection


Identity and isolation are evidenced; failure recovery and broader attack testing remain incomplete.



3/5 Medium confidence

Payments & entitlements


Forced write failure reproduced; no durable recovery.



2/5 High · blocker

Test & change safety

202 checks; revenue and AI recovery are not enforced.



2/5 High confidence

Release & rollback

Versioned process exists; failure proof is incomplete.



3/5 Medium confidence

Observability & response


Structured logs exist; payment drift is not auto-reconciled.



3/5 Medium confidence

Cost & provider operations

The £77 provider variance lacks complete attribution.



2/5 High confidence

1

Uncontrolled

2

Fragile

3

Managed

4

Resilient

5

Failure-tested

NO OVERALL SCORE

A launch blocker cannot be averaged away by strong documentation or architecture. The scoped launch decision remains separate and takes precedence over the scorecard.

SCOPE AND EVIDENCE

Four journeys, five evidence streams

The fictional review covered sign-in, subscription access, AI plan generation and PDF export across source, automated checks, staging runtime and read-only operational consoles.

JOURNEY	SOURCE REVIEW	AUTOMATED CHECKS	RUNTIME VALIDATION	PROVIDER CONSOLES
Sign in and session	Verified	Verified	Verified	Not applicable
Subscribe and unlock	Verified	Partial	Finding	Partial
Generate AI plan	Verified	Partial	Verified	Partial
Export plan	Verified	Verified	Verified	Not applicable

Each state applies only to the named illustrative evidence stream. A verified source review does not verify runtime behaviour.

SOURCE REVIEW

React/TypeScript client, server functions, schema, CI and operational documentation.

AUTOMATED CHECKS

164 unit/component, 34 integration and four end-to-end checks.

RUNTIME VALIDATION

Named staging build on desktop and mobile browser profiles.

PROVIDER CONSOLES

Read-only payment events, AI usage and crash/alert configuration.

HISTORICAL EVIDENCE

Two dated regression fixes were tied to their release records and passing checks.

EXCLUDED

Penetration test, load test, legal/compliance certification and production mutation.

HOW TO READ THIS

"Verified" means the named evidence supported the scoped check. "Partial" keeps a gap visible. It does not turn missing evidence into a pass or an automatic failure.

FINDING 01 · PAID ACCESS


A transient payment event can leave a paying user on the free plan

HIGH

Status: Open · 7 September 2026 Evidence confidence: High Launch blocker: Yes

OBSERVED IN FICTIONAL STAGING

The payment provider's event was acknowledged before the entitlement write completed. When the database write was forced to fail, the provider received success and did not retry. The user's payment remained visible in the provider console while the application continued to show free access.



BUSINESS CONSEQUENCE

Support must manually investigate a customer who has paid but received no value. Repeated cases create refunds, chargebacks and damaged trust at the conversion point.

LAUNCH DECISION

Blocker: yes. Fix before the fictional release candidate is promoted. This was reproduced safely in staging, not production.

Required outcome and verification


- Return a retryable failure until the entitlement grant is durable.
- Use a stable event key so retries cannot double-grant value.
- Add reconciliation for payment records that lack application entitlement.
- Run success, duplicate, delayed, out-of-order and forced-write-failure scenarios.
- Archive the test evidence against the exact release candidate.

RETEST RESULT REQUIRED

Pass only when the forced failure becomes visible, the provider retries, the grant completes once, and the user receives access without manual intervention.

MEDIUM

The application records successful plan-generation cost, but some retries, failures and missing telemetry remain only in the provider total.



Illustrative 30-day data · 10,912 successful plans

Ledger cost per success: $\text{£}351 \div 10,912 = \text{£}0.032$

Invoice cost per success: $\text{£}428 \div 10,912 = \text{£}0.039$

WHY IT MATTERS

Gross margin is being modelled from incomplete product telemetry. At small scale the cash difference is modest; during paid acquisition it can invalidate pricing and budget decisions.

EVIDENCE BOUNDARY

This sample assumes read-only product and provider totals for the same fictional 30-day period. The values are invented to demonstrate analysis.

Required outcome

- Record every provider attempt, including retries and terminal failures.
- Keep missing or unpriced units visible rather than treating them as zero.
- Reconcile product totals with provider billing on a schedule and alert on variance.
- Track cost per successful customer action, not only cost per request.

FINDING 03 · CHANGE SAFETY

Automated coverage is concentrated in happy paths

MEDIUM

Status: Open · 7 September 2026 Evidence confidence: High Launch blocker: No

The number of tests looks healthy. The journey matrix shows where assurance ends: subscription and AI recovery behaviour is not yet enforced in the release gate.

JOURNEY	HAPPY PATH	ALTERNATE PATH	UNHAPPY PATH	RECOVERY PATH
Sign in	Automated	Automated	Automated	Manual
Subscribe and unlock	Automated	Partial	None	None
Generate AI plan	Automated	Partial	None	Manual
Export plan	Automated	None	None	Not applicable

“Manual” means dated release evidence is required; “None” means the scenario is outside the enforced check set.

BUSINESS CONSEQUENCE

A release can pass 202 automated checks while the most expensive failure—payment without access—remains outside the enforced scenario set.

PROPORTIONATE RESPONSE

Do not chase a percentage or force every layer onto every function. Automate the stable, high-consequence scenarios and retain dated manual evidence for the rest.

Acceptance evidence

- Map each critical journey to happy, alternate, unhappy and recovery scenarios.
- Add request-level tests for payment-event failure, retry and idempotency.
- Add AI timeout, quota and retry-amplification checks.
- Fail CI if the selected high-risk suite is skipped or discovers zero tests.
- Record any deliberately manual scenario beside the release artifact.

AUDIT PRINCIPLE

Test presence is not assurance. The useful question is whether a risky change can reach users without a relevant check failing.

PRIORITIES AND HANDOVER

One blocker, then operating confidence

BEFORE LAUNCH	Repair subscription recovery and prove the exact release artifact. Owner: backend/release engineer. Depends on payment sandbox and staging failure injection.	1–2 days
BEFORE GROWTH	Close the AI invoice reconciliation gap. Owner: product/backend engineer. Depends on read-only provider billing and one aligned period.	1–3 days
BEFORE GROWTH	Enforce the risk-based journey matrix. Owner: release engineer. Depends on stable high-value scenarios and named manual device evidence.	1–2 days

Foundations already worth preserving

STRENGTH

Separate staging and production environments with synthetic accounts.

STRENGTH

Server-derived identity and policy-backed user-data isolation.

STRENGTH

Structured logs with correlation IDs and central sensitive-field redaction.

STRENGTH

Versioned setup, release, rollback and incident documentation.

WHAT THE PRIVATE REPORT ADDS

The complete paid deliverable adds every material finding with source/version evidence; the full A–L coverage disposition and test/gate inventory; explicit Not verified areas and unanswered questions; and a sequenced engineering handoff with owners, dependencies, effort bands and exact retest procedures.

Bounded retest

Recheck only the three findings against a named candidate. Mark each Pass, Fail or Not verified, then update the scoped launch decision. Do not imply that unchanged areas were re-audited.

DISCLOSURE

Asterpath Ltd is fictional. Every product detail, number, finding, test result and remediation estimate in this sample is illustrative, and any resemblance to a real business is coincidental. The document shows what a gadirlabs deliverable can look like; it is not proof of a client engagement or a claim about any real company.

A real audit would use the client's agreed scope, repository and runtime evidence, read-only consoles, explicit limitations and an owner-approved publication boundary.