

Technical Feasibility Review

Expert marketplace platform · v2.0 · Redacted sample

Redacted sample. This is an example of a real report, prepared for a founder on a real product idea and edited here for external sharing. The company and all named individuals have been removed, and a small number of figures have been widened into ranges. The findings, the delivery plan and the confidence ratings are unchanged.

Report details

Field	Detail
Prepared for	Withheld
Prepared by	Stephen Goggin, Fractional CTO
Engagement	Technical Feasibility Review (five working days)
Version	2.0
Basis of review	Two working sessions with the founder, a product requirements document, two versions of a user interface specification, a click-through demonstration of ten screens, a market research document, a validation plan, and for context only a competitor analysis and an outreach contact list. No source code was provided, as none exists.
Confidentiality	Confidential to the client. Suitable for release to investors and advisers.

Revision history

Version	Date	Note
1.0	Withheld	Initial issue
2.0	Withheld	Redacted for external sharing

Intended use

Suitable for	Not suitable for
Founder decision-making	Security certification or assurance
Investor and board discussion	Legal advice
Funding and grant applications	Financial or commercial advice
Technical planning and budgeting	Commercial validation of the proposition

Estimates. All costs and durations in this report are directional planning estimates produced within a time-boxed review. They are not quotations, tenders or commitments, and they are not a substitute for a supplier quotation or a detailed estimate produced against a settled specification.

Questions this report answers

The engagement was scoped around the questions below, agreed at kick-off. The client's own primary question is answered first.

Question	Answer in brief	Section
What will it cost to build, and how soon is something in front of users?	€95,000 to €135,000 for the build. Verified experts and working discovery after ten to twelve weeks of build, excluding any hiring lead time. Whether that sum can be raised is a commercial question outside this review.	1, 4
Can the product be built as described?	Feasible, subject to changes of scope and sequence.	1, 3
What is genuinely difficult about it?	Holding funds between booking and delivery, selective anonymity, and verification.	3.2
What should be bought instead of built?	Payments, identity verification, search, messaging, email, storage and monitoring.	3.4
What should come out of the initial release?	Five items, each with the condition that brings it back.	3.7
What will it cost to run, and who has to run it?	€85 to €395 a month at launch, and four operational functions that carry high human effort.	3.8, 4
Who is needed to build it?	Two senior engineers and one non-engineering reviewer, with part-time product and visual design. No scarce skill is required.	3.10
What could delay it?	Six causes, with the delay quantified where it can be.	1
What regulation affects the build?	Eleven areas named, with the one to address first.	6

Questions this report does not answer

- Whether anyone wants the product, or will pay for it.
- What to charge, for anything.

- Whether the competitive position is defensible.
- Whether the team can execute.
- Whether the security implementation is sound. Security was not assessed; see section 8.
- The legal position in any regulatory area named. Those are identified and characterised for planning, not advised upon.

What this report enables

- Approaching a supplier or contractor for a quotation against the phasing in section 4.
- Briefing an engineer or an agency on the approach from sections 2, 3 and 4. A specification complete enough to hand over for delivery is further work, as section 4 records.
- Attaching to a grant or funding application as technical evidence.
- Using section 7 in investor materials without rewriting it.
- Commissioning the phase 1 investigations, which are bounded, priced and independent of the rest.
- Deciding whether to proceed, alongside the commercial information this review did not see.

Scope of this assessment

This is a **technical** feasibility review. It addresses whether the solution described can be built, the components it would be built from, the indicative cost and duration, and the technical risks that apply.

It is not a market assessment. This review takes no position on demand, willingness to pay, pricing or market size. A finding of technical feasibility confirms that the engineering is achievable. It does not indicate that the engineering is commercially justified.

Feasible, in this report, means three things: the solution requires no unsolved technical problem; it can be delivered with technology available today from suppliers that exist; and it can be delivered within the assumptions recorded in appendix A. A finding of feasible carries no view on whether the solution should be built.

Assessment outcomes

One of the following is recorded in section 1. The definitions are fixed so that reports remain comparable across engagements.

Outcome	Meaning
Feasible	No material technical barrier identified.
Feasible with constraints	Technically achievable, subject to identified assumptions, investigations or changes of design or scope.
Not feasible as described	The proposed design cannot meet its stated objectives without substantial redesign.
Insufficient evidence	Further investigation is required before a conclusion can be drawn.

Within scope	Outside scope
Can the solution be built as described?	Is there demand for it?

Within scope	Outside scope
Indicative cost and duration	Willingness to pay, and at what price
Principal technical challenges	Market size
Build versus buy recommendations	Competitive position
Technical risk exposure	Team capability
Regulatory impact on the build	Route to market

Confidence ratings

Each material judgement in this report carries a confidence rating. A one-week review cannot establish everything to the same standard, and the ratings identify where the evidence is strong and where it is not.

Rating	Basis	To raise it
High	Verified directly during the review, or based on well-established implementation patterns.	Further investigation is unlikely to change the conclusion.
Medium	Established approach, dependent on project-specific assumptions not yet confirmed.	A technical investigation of one to three days.
Low	Requires validation through investigation, proof-of-concept work or vendor confirmation.	A proof of concept or written vendor confirmation before material expenditure.

1. Executive summary

Scope: technical feasibility only. This assessment does not address demand, pricing or commercial viability.

This section is written to be read alone. It carries the conclusion, the cost, the schedule, the staffing required and the risks to each, so that it can be weighed against commercial and financial information the review did not see. The sections that follow set out the evidence behind it.

Evidence reviewed

Two working sessions with the founder, a product requirements document, two versions of a user interface specification, a ten-screen click-through demonstration, a market research document, a validation plan, and for context only a competitor analysis and an outreach contact list. No source code exists. No vendor was contacted and no contract was reviewed. Appendix B lists the material in full and appendix C records what was relied upon without verification.

The proposed solution

A marketplace enabling a buyer requiring subject-matter expertise to search verified experts, book time and pay through the platform, without either party exchanging contact details. Experts are verified before they are published, and each controls when, up to the point of booking, their identity is released. The platform intermediates discovery, scheduling, payment and reviews. The estimates in this report assume the platform retains a share of each transaction, which is the model the requirements describe. Section 4 records that this has not in fact been decided, and that the alternative, a listing fee, is a different system rather than a different price. Phase 4 would be rebuilt around recurring billing instead of held funds and payout, the phase cost would change accordingly, and expert registration in phase 2 would collect different data. Launch is into Ireland and the wider EU, English-only, with the operating company to be established in Ireland and the platform EU-hosted.

Technical feasibility assessment

Feasible with constraints. Achievable against the technology, subject to the investigations and scope changes set out below. Delivery is further constrained by budget, schedule and staffing, which this review sizes but does not decide.

Basis: the solution requires no novel technical capability. Verified supply, scheduled sessions and intermediated payment are well-established patterns supported by mature commercial components, and most competent teams would reach the same technical conclusion. That conclusion is not what decides the project. What decides whether it can be delivered is whether the budget, the elapsed time and the two senior people set out below can be secured. Whether it should proceed is a separate, commercial judgement, addressed at the end of this section. The recommended changes to the product itself concern sequence and scope: five elements should be deferred from the initial release (section 3.7), and one payment dependency resolved before the booking flow is designed.

The binding constraints

Constraint	Position	What it means for the decision
Budget	€95,000 to €135,000 for the build alone	Excludes product and visual design, content, legal, the verification reviewer's salary, an external security review before launch, marketing and support. Cost to launch is materially higher and is not established by this review.
Elapsed time	Five to seven months in-house; six to	The build is 19 to 22 weeks. Phase 1 runs while the

Constraint	Position	What it means for the decision
	nine months outsourced	team is hired, so the in-house figure is the four-to-ten-week hiring lead time plus phases 2 to 5. The outsourced figure adds a vendor lead time of two to four months to the build. See schedule risk below.
Staffing	Two senior engineers, plus one non-engineering reviewer	No rare or specialist technology skill is required. The constraint is securing people at that level, and how quickly. This is not a headcount that juniors can fill. See section 3.10.

Indicative delivery estimate

Stage	Duration	Indicative cost	Confidence
Technical investigations	1 week	€5k – €6k	High
Early release, introductions handled manually	10 – 12 weeks	€50k – €74k	Medium
Soft launch, self-serve and paid	19 – 22 weeks	€95k – €135k	Medium

Durations are weeks of build, with phase 1 followed directly by phase 2. Where phase 2 waits for the team to be hired, the elapsed times under the binding constraints apply instead.

Operating costs: €85 to €395 per month at launch, rising to €800 to €2,000 at twelve months (section 4).

Schedule risk

The durations above assume the following hold. Each of these has delayed comparable projects, and the effect is given in elapsed time.

Cause	Effect on schedule	How to remove it
Payment provider does not confirm in writing within the investigation window	2 to 4 weeks, and phase 4 expands if the answer is unfavourable	Contact candidate providers as soon as the entity type is settled, before phase 1 formally begins
Specification continues to change after phase 2 starts	Pushes cost and duration to the top of both ranges	Freeze the specification, having first resolved the divergence in section 9
Legal opinions outstanding on erasure, held funds or invoicing	Phase 2 or phase 4 cannot start, depending on the opinion; section 6 sets out which	Instruct the erasure opinion in parallel with phase 1, the held-funds opinion immediately if the phase 1 confirmation is adverse and otherwise before phase 4, and the invoicing opinion early in phase 2
Verification policy for non-credentialed experts undecided	Phase 2 cannot close	Decide during phase 1, while the manual check is being timed
Senior engineers not secured	4 to 10 weeks before phase 2 starts	Begin hiring immediately. Phase 1 runs in the meantime and does not depend on the team

Cause	Effect on schedule	How to remove it
		being in place
Outsourced delivery selected	2 to 4 months before development begins	Decide the delivery model now; see section 4

Principal risks

- **The booking flow as specified may not work.** It relies on a card authorisation lasting from booking to the session, which the standard window does not cover. Extended authorisation may, subject to provider support and commercial conditions. The design recommended in section 3.2 avoids that dependency by capturing at booking, which in turn requires the provider to permit funds to be held. Confirm before the booking flow is designed.
- **Verification is an operational cost.** No single automated credential check covers the launch market. A share of registrations can be confirmed immediately; the remainder carry a manual cost that recurs. The proportion, and therefore the cost, must be established before the verification fee is set; this review does not assess the fee.
- **Identity disclosure through secondary channels.** Calendar invitations, receipts and search filters all carry identifying information by default. Each is a small fix; one omission defeats the core proposition.

Critical dependency affecting this assessment. If funds cannot be held between booking and delivery for the entity type to be registered, the booking model requires redesign and phase 4 expands. The first investigation is intended to resolve this within three days, provided the entity type is settled and the provider confirms in writing.

Recommended immediate actions

Two technical investigations, totalling one week for one senior person, before further commitment. They can begin as soon as the entity type is settled, carried out by the lead engineer if already engaged or otherwise by an adviser, while the rest of the team is hired. Up to three days to select a payment provider and obtain written confirmation of how funds may be held between booking and delivery; the confirmation itself is half a day and the remainder covers provider selection. The confirmation depends on the entity type, which is a company decision and should be settled before phase 1 starts. Two days to measure the reviewer effort involved in one manual verification of a non-credentialed expert; any external reference requested during that work is recorded as a turnaround figure rather than awaited. These results determine the shape of the build.

Critical assumptions

The assessment rests on the following. Each is recorded in full in appendix A, except the scale targets, which are recorded in appendix C; they appear here because a reader should not have to reach an appendix to find what the estimates and the recommendations depend on.

Assumption	Status	Consequence if incorrect
Funds can be held between booking and delivery under a provider's authorisation	Unverified	Booking model redesigned and phase 4 expands; the payment design and timeline change, though the feasibility conclusion does not
Users on both sides, and the controller,	Stated by the client,	The regulatory position in section 6

Assumption	Status	Consequence if incorrect
are within the EEA	unverified	changes and must be reassessed
The demonstration, not the specification, is the authoritative artefact	Unconfirmed	Parts of sections 2 and 3 were assessed against the wrong material
Verification can be tiered, with the credential check applied selectively	Proposed, untested	Cost per registration and operational load both rise
Scale targets as supplied by the client	Unsupported	Infrastructure sizing and operational staffing change

What this review settles

The following no longer require deliberation unless the assumptions in appendix A change: that the product can be built without novel technology; that the architecture should be conventional, with the payment provider held behind an internal interface; that five features are deferred from the initial release, including permanent anonymity, whose data model requirement is settled since identity must be held separately in any case, though whether the feature itself is ever built awaits the legal advice recorded in section 4; that credential verification cannot be treated as a single automated step in the launch market; and that no scarce technical skill is required to build it. The tiering proposed in section 3.5 is a response to the credential-verification finding and is itself untested, so it is recorded as an assumption rather than as settled. Sections 4 and 6 record the decisions that remain open and the phase each one gates.

What this assessment does not decide

Whether to proceed is a commercial judgement and this review does not make it. The figures above are one input. The others, none of which were within scope, are whether there is demand, what buyers will pay, what the competitive position is, and whether the funds and the runway exist to carry the five to nine months from decision to launch before revenue. A decision requires this document alongside those.

2. Solution overview

Overview

The transaction runs in five steps. A buyer searches and shortlists. The platform carries an enquiry between the parties without exposing contact details. The expert accepts and confirms a price. A session is scheduled and the buyer pays, with funds held until delivery. The session takes place, the expert is paid less commission, and both parties may leave a review. Identity is released at a point of the expert's choosing somewhere in that sequence, no later than booking in the initial release, and that single variable shapes most of the build.

Markets and establishment

Launch is into Ireland and the wider EU. Both sides of the marketplace are expected to be within the EEA at launch, and the operating company is to be established in Ireland. English only at launch, with translation deferred. These points are recorded here because most of section 6 depends on them: the regulatory position follows from where the controller is established and where the users are, not from where the servers sit. Expansion beyond the EEA on either side brings further regimes into scope and is not assessed here.

Scale assumed

The figures used throughout this report are the client's own targets: 500 experts, 200 individual buyers and 5 corporate accounts at six months, rising to 10,000 experts at twenty-four months. The twelve-month operating cost is estimated between those two points. The targets were supplied by the client, are not supported by evidence this review examined, and are recorded in appendix C. They are used to size infrastructure, to estimate operating cost and staffing, and to test whether any component requires unusual engineering. The feasibility conclusion does not depend on them. The operating cost in section 3.4 and section 4, the operational staffing in section 3.8, and the recommendation to defer availability-filtered search do, and each would need revisiting if the targets prove materially wrong.

Target users

Supply side: individuals holding recognised subject expertise who receive irregular requests for their time and have no structured means of handling them. Typical profiles include academics, former regulators and established practitioners in narrow fields. Demand side: professionals working to deadline who require a credible expert voice at short notice, currently relying on personal networks, general search or repeat use of previous contacts.

Intended business outcome

For the expert, discoverability without publication of personal contact details, and payment without direct invoicing of an unknown counterparty. For the buyer, a shortlist of verified individuals with stated availability and price, in place of an unqualified list of names.

Core technical assumption

That an expert's identity can be withheld throughout discovery and released at a point of the expert's choosing, no later than booking in the initial release, while payment, calendar invitations and messaging continue to operate correctly around that constraint. The distinctive characteristics of the solution depend on this assumption holding. The assessment finds it achievable; section 3.2 sets out what implementing it properly requires.

3. Technical assessment

3.1 Proposed architecture

A conventional web application: a server-rendered front end, an application programming interface, a relational database as the system of record, object storage for documents and images, a separate search index, and the third-party services set out in section 3.4, of which payments and payouts, identity verification, transactional email and real-time messaging are the ones that shape the design.

The defining characteristic of the design is that several subsystems converge on a single event. On confirmation of a booking, the expert's identity is released, a calendar invitation is issued and funds are held; once availability filtering is introduced (section 3.2), the search index is also updated to reflect the allocated slot. That event concentrates the complexity of the solution; the components either side of it are routine. The booking should therefore be implemented as an explicit state machine with recorded transitions from the outset. States inferred from record-level flags will not survive the diagnostic work this component attracts over the platform's life.

Why this shape. Delivery risk is the criterion the architecture is chosen against. Every component is a pattern with mature library and vendor support, which keeps the hiring pool wide and the failure modes documented. State is held in a single relational database so that the booking, the funds and the disclosure decision cannot diverge from one another. Third-party services are used wherever the problem is already solved commercially, which moves cost from build to subscription and shortens the initial release. The phases can each be delivered to a working system, which is what makes the review point in section 4 possible. Nothing in the design anticipates volumes beyond those recorded in section 2.

Technology selection. A specific technology stack is named in the project documentation. It was not agreed with the founder and is not treated by this review as a settled decision. Any mainstream server-side language with a mature payments library and a well-supported relational database is suitable, and selection should follow the experience of the initial engineering team. The material decision is to abstract the payment provider behind an internal interface so that an additional provider can be introduced later. This requirement was identified by the client and is retained in the recommended architecture.

3.2 Technical challenges and design considerations

Holding funds between booking and delivery

This is the highest-priority dependency identified by the review. The flow in the requirements authorises the buyer's card at booking and captures the payment after the session. A card authorisation is a temporary hold and expires. The standard window is approximately seven days, varying by card scheme and acquirer, against an assumed booking lead time of two to four weeks. Extended authorisation lengthens the window to approximately thirty days on the major schemes, but only where the acquirer supports it, on interchange-plus pricing, and in some cases for a per-transaction fee. The flow as specified therefore works only if extended authorisation is available to this platform, which has not been established.

Three designs are available. The first captures payment at booking and holds the funds within the provider's balance, releasing payment to the expert after delivery. The second authorises closer to the session date, accepting that a proportion of charges will fail at the least convenient moment. The third takes a deposit at booking and the balance on the day. The first is recommended. It is the intended use of the connected-account products offered by the major providers, it supports a clear cancellation policy, and it removes the dependence on the authorisation window. The dependency moves rather than disappears: the provider must permit funds to be held between booking and delivery for the entity type

to be registered, and client funds are then held for a period, under the provider's authorisation where its model covers that activity; section 6 sets out what follows if it does not.

Strong Customer Authentication bears on the choice. Published regulatory guidance places a charge made without the cardholder present outside the authentication requirement, and attaches the requirement instead to the point at which the card is stored and the payment mandate agreed. The second design depends on exactly such a later charge. It succeeds where the mandate was taken and authenticated correctly at booking and fails where it was not, and because the cardholder is absent the issuer declines rather than challenges, so the payment cannot be recovered in the moment. The first design keeps the cardholder present for authentication, which is a further argument for it. The treatment should be confirmed with the provider selected, and appendix C records it as unverified.

Two related items were not examined and should be settled with the provider before the payment design is fixed: how chargeback liability is allocated between platform and expert under a connected-account model, and which PCI self-assessment level applies to the integration chosen.

Confidence: Low. The mechanism is standard, but whether the provider permits funds to be held for this entity type, and whether extended authorisation is available, both await vendor confirmation, which the rating scale places at Low. Confirmation is half a day of the phase 1 investigation set out in section 1, and depends on the entity type being settled first.

Selective anonymity

Selective anonymity constrains the data model. It cannot be switched on later. Where an expert is anonymous during discovery and identified subsequently, identifying attributes must be held separately from the profile under distinct access controls, and every surface touching a booking must be assessed for disclosure. The principal surfaces present no difficulty. The disclosure risks most often overlooked arise in routine integrations: calendar invitations carry the organiser's name and address by default, payment receipts identify the payee, email headers expose the sending domain, and uploaded files retain embedded metadata.

A secondary risk is re-identification through filter combinations. Where a search is narrowed by specialism, jurisdiction and seniority, and few profiles match, anonymity may be defeated without any misuse of the system. The established mitigation is to suppress or generalise results where a filter combination returns fewer than a defined minimum number of anonymous profiles. The technique is well understood and inexpensive, but must be designed in from the start.

Confidence: Medium. The approach is established. The number of disclosure points depends on the vendors ultimately selected, each of which will need individual assessment.

Invoicing and VAT against anonymity

A booking between an EEA platform, an EEA expert and an EEA buyer is a VAT-relevant supply, and where the buyer is a business it generates an invoice identifying the supplier. That is the information selective anonymity withholds, and the report identifies the collision without resolving it. Three routes exist. The platform invoices as principal and the expert invoices the platform, which resolves the conflict but moves the platform from intermediary to principal, with consequences for VAT treatment and for the commission model. The platform issues the invoice in the expert's name and on the expert's behalf, which is third-party invoicing rather than self-billing, and which still discloses the expert to the buyer. Or anonymity is released at the point of invoicing, which confines anonymity to the discovery stage. The third sits closest to the disclosure-by-booking model recommended in section 3.7 and should be tested first.

Place-of-supply treatment, VAT registration number capture and validation, and the invoicing route should be settled before phase 4, as each produces a different data model. Confidence: Low. The conflict

was identified during the review and not researched to a conclusion. It requires investigation and advice, and section 6 records it as an area requiring advice.

Two verification populations

Verification of an academic against an institutional register is, in principle, an automated lookup. Verification of a former regulator, a practising clinician or a retired commercial pilot is not, as no equivalent register exists. Assessed against the expert categories set out in the project documentation, most of the intended supply falls into the second group.

For the first group, section 3.5 records that automated confirmation is available only for some institutions in the launch market. For the second there is no register to consult at all. Section 3.5 sets out a tiered response. The identity check in the base tier is automated; corroboration for the second group needs a reviewer. The credential tier is an operational process requiring human review wherever no automated route exists, and automating it further should not be attempted in the initial release. The recommended implementation is a workflow with explicit states, a documented evidence trail, a reviewer queue and a record of each approval decision and its author. The development effort is modest. The material cost is reviewer time, and it recurs, as a credential with fixed validity returns the entire population for re-verification.

Confidence: Medium. The absence of a register for these categories is established. The proportion of the intended supply that falls into the group has not been tested against a recruiter in each field, and appendix C records it as unverified.

Availability-aware search

Filtering search results by availability costs more than it appears. Availability must be held within the search index to be filterable, requiring the index to be updated on every booking, cancellation and calendar amendment. This constitutes a denormalised copy of live data, which is subject to drift and is a recognised source of defects in which the system offers a slot that is no longer available.

Availability filtering should be deferred from the initial release. A general indication on the profile, with a definitive check at the point of booking, is sufficient. Filtering can be introduced once supply density makes it meaningful. At the launch volumes assumed, the feature delivers limited value against two weeks of development effort.

Confidence: Medium. The engineering cost is well established. The judgement of value rests on the client's scale targets, which appendix C records as unsupported, and would need revisiting if they prove materially wrong.

3.3 Areas of lower technical risk

Scheduling implemented in-house. A booking calendar is to be built in-house, and that decision will attract challenge. It is the right one here. Availability rules, time zone handling and invitation issuance are well-understood implementation areas with proven solutions available. Bi-directional synchronisation with a user's work calendar, reading existing commitments and writing back, is a different order of problem. It was explicitly excluded. The exclusion is retained in the recommended scope. Times must be stored as UTC alongside the named time zone rather than as a fixed offset, as offset-based storage produces a recurring class of defects at daylight saving transitions.

Messaging. Hosted services and established open-source servers both cover real-time messaging between two parties. Building a messaging layer from first principles is not recommended.

Scale. At the volumes projected for the first two years, the data set is small. No unusual infrastructure is required. Proposals introducing it should be tested against projected volumes before acceptance.

Translation. Deferred beyond the initial release. When introduced, it constitutes an application programming interface integration with appropriate labelling of output, and is not a large work package.

3.4 Recommended third-party services

The costs below are published list prices as at the date of review. No quotations were obtained and several are volume-dependent, so these indicate the scale of expenditure only.

Requirement	Recommended approach	Indicative cost
Application hosting and database	Managed application platform and managed PostgreSQL in an EU region, with daily backups and a small cache. Sized for launch volumes; no redundancy beyond the provider default. The lower figure reflects published list prices at a mid-market provider, the upper a managed cloud build where the load balancer and gateway carry a fixed monthly charge before any traffic.	€40 – €150 per month at launch
Payments and payouts	Stripe Connect, abstracted behind an internal interface. Its current account model should be used; the older Express account type now carries a deprecation notice for new platforms. Adyen for Platforms and Mangopay are alternatives. They are not equivalent on the question that matters here: Mangopay is a licensed e-money institution whose model holds funds per user natively, whereas Stripe constructs comparable behaviour through delayed payouts. That distinction bears directly on the first investigation in section 4.	Approximately 1.5% + €0.25 on standard EEA consumer cards; commercial and corporate cards are materially higher, which matters given the corporate demand side. A per-connected-account monthly charge and per-payout fees apply in addition and scale with the number of active experts.
Identity verification	Veriff, or Stripe Identity where Stripe is already the payment provider. Entrust, which acquired Onfido, is a further option at enterprise scale.	Roughly €1 – €2 per check on published self-serve tiers. Where the vendor applies a monthly minimum, roughly €45 at the entry tier, the minimum rather than the per-check price sets the cost at launch volume
Authentication	Auth0 or Clerk, or a mature framework library where the team prefers to self-host	Negligible at projected scale
Search	Typesense, self-hosted within an EU region. Elasticsearch is heavier than required here; Algolia was excluded on cost, not on data residency, since	Infrastructure only, €30 – €80 per month at launch

Requirement	Recommended approach	Indicative cost
	it operates EU regions.	
Transactional email	Postmark	€15 – €50 per month at launch
Real-time messaging	Abyl, or a self-managed Socket.io deployment	€0 – €50 per month at launch
File and image storage	Cloudflare R2, or AWS S3 in eu-west-1	Under €20 per month at launch
Error and performance monitoring	Sentry	Free tier sufficient initially
Translation (deferred)	DeepL, with translated output labelled as such	Per character, negligible until volume exists

No vendor has been approached, priced or contracted. Those named are the candidates assessed against this solution, with pricing taken from their published pages during the engagement. The payment provider is selected within phase 1; final selection of the remaining vendors requires about a day once the entity type is confirmed, and several of the selections are constrained by the data residency decision recorded in section 6.

3.5 Data and integrations

The platform holds expert profile and credential data, message history, booking records, payment records and review content, together with the outcome of each identity check. The documentary evidence behind an identity check should not be retained by the platform at all; the verification outcome, timestamp and provider reference are sufficient, with the evidence held by the provider. The most sensitive holding is therefore the segregated store of anonymous experts' identities, which requires access control and audit logging separate from general administrative access.

The solution has no inbound integrations with third-party systems, which removes a whole category of delivery risk. All external dependencies are commercially procured services.

Credential verification in the launch market

The design assumes a single lookup service against which academic credentials can be checked. No single service covers the launch market, and the position varies by institution. HEDD, the service most commonly cited for this purpose, covers United Kingdom institutions only. A shared digital credential platform operates across a substantial number of Irish higher education institutions: a graduate shares a verifiable record directly, confirmation is immediate, and an interface is offered for volume use. QQI operates an employer verification service for awards made by QQI and its predecessor bodies, at €20 per check with a stated turnaround of up to fifteen working days. Awards from institutions outside the shared platform are verified individually through each institution's own registry service, on separate terms and timescales.

The consequence for the design is that credential turnaround is bimodal: immediate for part of the population and measured in weeks for the rest. Where an award falls inside the shared platform's coverage, the check can be built into registration. Elsewhere it is a manual request with an external service level that no development effort shortens.

The design response is to tier verification instead of applying one standard across the population. A base tier confirms identity through the identity verification service in section 3.4 at registration, ahead of the payment provider's own checks when payouts are set up in phase 4, and gathers corroborating signals for

the expert's stated field: a confirmed institutional or corporate email domain, a public professional profile, a publication or media record, or membership of a professional register. The identity check within it is automated and inexpensive. Corroborating an expert outside any register needs a reviewer to weigh the signals, which is the effort measured in phase 1 and, on the categories in the project documentation, the majority case. A credential tier adds the formal registry check and applies only where an expert claims a specific academic award and wants it displayed, or where a buyer requires that assurance. The tier then becomes an attribute on the profile, so the buyer selects the level of assurance and platform cost scales with what is actually asked for.

Identity, credentials and experience are three separate claims

Only one of the three can be checked cheaply, and it is not the one the buyer is paying for. An identity check confirms that a person is who they say they are, which is what such checks exist to do. They establish nothing about whether that person is an immunologist, spent ten years at a regulator, or holds the experience a buyer is engaging them for. A registry check extends to a named award and no further. Employment history and practical experience are established by contacting the employer, which is a reference process: manual, slow, and dependent on a third party choosing to respond. A public professional profile is corroboration rather than verification; it evidences a consistent public presence, not the truth of what it asserts.

The design consequence is that the profile must state which claims were checked and by what means, and must not present an identity check as though it covered the rest. Three separate indicators can be displayed without overstatement: identity confirmed, credential confirmed, and experience self-declared. A single verified badge covering all three presents the platform as having checked more than it did. Whether that creates liability is a question for the company's adviser; the design point is that the claim should not be made. This is the same principle as section 3.6: publish the components, not a composite that implies more than was established.

A profile should be created in a pending state and published once the base tier completes. The credential indicator is added when the formal check returns, which for part of the population will be weeks after registration.

Confidence: Medium on the coverage position. The QQI and HEDD terms were checked against their published pages; the institutional coverage of the shared credential platform was established from its own published list and was not confirmed institution by institution, and neither its commercial terms at volume nor its interface were tested. Low on per-institution turnaround, which varies by university and was not established for each; raising it requires a written service level from each institution.

3.6 Artificial intelligence and machine learning

The solution has limited dependency on machine learning. Translation is a commercial application programming interface. Detection of attempts to move a conversation off-platform is pattern matching with a human review queue. No model is required. Flagging for review in place of a real-time block is the correct choice: a false positive that blocks a legitimate message does more damage than one that queues it.

The credibility score requires reconsideration. The design proposes a single score out of one hundred, derived from weighted components, recalculated nightly and presented in bands. The weights recorded in the documentation are not derived from data, as no outcome data yet exists against which to calibrate them. A composite score constructed from unvalidated weights carries an appearance of objectivity it cannot support, and would be difficult to defend should an expert challenge their relative ranking. The recommended approach is to publish the underlying indicators directly: the three verification indicators set out in section 3.5, completed engagement count, review average with sample size, and date of last

activity. These are individually defensible and address the buyer's information requirement. A composite measure can be introduced once sufficient completed bookings exist to establish which indicators predict a satisfactory outcome. Confidence: High that the current design is premature. Low as to whether a useful composite is achievable, which cannot be tested until outcome data exists.

3.7 Recommended scope reductions for the initial release

The items below carry the highest development cost relative to the validation they deliver at this stage. Their deferral accounts for most of the variance against the earlier internal estimate.

Defer	Rationale	Reintroduction trigger
Availability filtering in search	Requires live calendar data within the search index, maintained against every booking and cancellation. Approximately two weeks of effort and a persistent source of stale results.	Sufficient supply density for filtering to be meaningful.
Permanent anonymity	Identity is never disclosed, which leaves open questions of contract formation and liability that this review does not resolve. Requiring disclosure no later than booking means both parties are identified when the commitment is formed, which avoids the anonymous-counterparty situation those questions concern, while anonymity is retained through discovery.	Receipt of a supportive legal opinion.
The credibility score	A composite derived from unvalidated weights. The underlying indicators are more defensible.	Sufficient completed bookings to calibrate the measure.
Translation	A contained integration with limited value at launch.	Demonstrated demand outside English.
Advertising on the free tier	Introduces a consent mechanism and a second revenue system, both of which carry build and operational cost from the initial release. This review does not assess the return, which is a commercial question for the company.	A decision by the company that free-tier traffic justifies the additional systems.

3.8 Operational load

The design implies continuing human work, and that work is a running cost. It is set out here because it is routinely underestimated and because it determines staffing before revenue.

Function	Software effort	Human effort	Note
Verification, base tier	Medium	Medium	Identity checks are automated.

Function	Software effort	Human effort	Note
			Corroborating experts outside any register needs a reviewer, which is most of the intended supply.
Verification, credential tier	Low	High	Manual outside the shared credential platform, and recurs on re-verification.
Review dispute handling	Low	High	The expert may contest before publication. No written policy exists.
Refunds and cancellations	Medium	Medium	Policy-driven. Most can be automated once the policy is set.
Abuse and off-platform reports	Low	High	The flag-and-queue design is deliberate and puts a person in every decision.
Identity and anonymity complaints	Low	High	Low volume, high consequence. Needs a named owner.

Four of the six functions carry high human effort against low software effort, so the operating model determines whether they work. Written policies are required for review disputes, refunds and abuse handling before phase 4, because each changes what the software must record. Confidence: Medium. The volumes behind these are unknown, and appendix C records the scale targets as unsupported.

3.9 Deferred scope and expected rework

The scope reductions in 3.7 are deliberate and they are not free. Recording where rework is expected is more useful than presenting the initial release as a finished foundation.

Area	Expected rework
Search	Availability filtering requires a denormalised slot index and an update path on every booking event. That is an addition to the search architecture, not a configuration change.
Verification	Moving any part of the credential tier towards automation requires the workflow to carry evidence in structured form from the start. Designing for that now costs little; retrofitting it does not.
Reporting	Seller income reporting, and any board or investor reporting, will need data the initial release is not otherwise obliged to retain. The retention decision must be settled before phase 2 and implemented in it.
Anonymity	If permanent anonymity is later brought into scope the data model must already separate identity. That separation is in the recommended design; the interface and policy work would be new.

3.10 Skills and availability

Nothing in this solution requires a rare or specialist technology skill. The stack is mainstream, the payment and verification vendors publish documented interfaces, and the patterns are common to any marketplace. That finding is useful in itself: the delivery constraint is the speed of securing people.

Role	Availability	Note
Senior engineer, payments and state machines	Good	Widely available in the Irish contract market. Experience of connected-account payouts narrows the field but is not rare.
Senior engineer, data model and privacy	Good	The requirement is seniority. This role holds the anonymity model and the disclosure audit, and should not rotate.
Front-end and interface	Good	No unusual requirement. Covered by the two senior engineers in the estimate, and suitable for outsourcing once the design is settled if capacity is short.
Verification reviewer	Good, but non-technical	Not an engineering hire. Part-time at launch volumes, and recurring. Frequently omitted from startup plans.
Product and visual design	Good	Part-time. Excluded from the cost figures, as section 4 records.
Security review before launch	Good, external	A bought engagement. Budget for it separately.

The privacy and data model role is the single point of concentration in this build; a team that loses it mid-project loses the reasoning behind the disclosure design. The verification reviewer is an operating cost that begins before revenue. Confidence: Medium on availability, which reflects the Irish contract market at the date of review.

4. Delivery plan and indicative costs

Phased so that each stage resolves a defined question. Durations assume one senior person for phase 1 and two full-time engineers thereafter, with part-time product and visual design and verification administration. Costs are stated as ranges, at Irish contract rates for phases 2 to 5 and a senior advisory rate for phase 1, and exclude VAT, product and visual design, legal and content.

Phase	What it resolves	Duration	Indicative cost
1 — investigations	How funds are held between booking and delivery, and the effort required for one manual verification. One senior person.	1 week	€5k – €6k
2 — supply side	Registration, profile, credential evidence, the verification workflow and administrative queue. Concludes with verified experts published.	5 – 6 weeks	€25k – €37k
3 — discovery and enquiry	Search, profiles under enforced anonymity rules, and intermediated enquiry. No calendar or payment.	4 – 5 weeks	€20k – €31k

Phase	What it resolves	Duration	Indicative cost
	Introductions completed manually.		
4 — booking and payment	Calendar, confirmed bookings, funds held, payout after delivery, reviews including the dispute stage.	6 – 7 weeks	€30k – €43k
5 — hardening	Disclosure audit across all surfaces, load verification, accessibility review, administrative tooling.	3 weeks	€15k – €18k

Basis of the ranges. Contract rate is the principal variable. The lower bound assumes a settled specification; the upper bound allows for the level of specification change usual on a first product.

Delivery model

The scope is 185 to 215 person-days. Phases 2 to 5 account for 180 to 210 of them, for two engineers, costed within a band of €500 to €620 per day. Published Irish contract guidance for 2026 places software engineers with five or more years at roughly €395 to €505 a day and lead, principal or staff engineers at roughly €505 to €620. The band used runs from near the top of the experienced range to the top of the lead range. Its upper figure prices both engineers at the top of the lead range, which is deliberately conservative. A rate above €620 for that work is outside the published Irish contract range and should be questioned. The technical design is carried out by the same engineers, mainly at the start of phase 2, and is included in the phase figures. Phase 1 is five days for one senior person, carried out while the rest of the team is hired, and is priced at €1,000 to €1,200 a day, within the published European range for senior advisory work, so that the figure is an upper bound whoever carries it out; if the lead engineer does it at contract rates, phase 1 costs less. Because it needs one person and no team, it can be commissioned on its own.

An outsourced or offshore team is the obvious alternative. Published agency rates in 2026 run at roughly \$40 to \$100 per hour in Latin America and \$25 to \$65 per hour in South and South-East Asia. Applied to the 180 to 210 person-days of phases 2 to 5, at eight hours a day and an exchange rate of about 0.91, that is a spread of roughly €33,000 to €153,000, against an in-house figure of €90,000 to €129,000 for the same phases. The lower part of that spread undercuts the in-house figure substantially; the top of the Latin American band exceeds it. Any saving is therefore contingent on where in the band an engagement lands and on the vendor holding that rate for the duration. Confidence: Medium on the rate comparison, which uses published agency ranges rather than quotations.

The difficulty in this solution sits in judgement. The disclosure surfaces in section 3.2, the payment hold design, the verification policy and the regulatory constraints in section 6 each require a decision, and each has to be specified precisely before it can be delegated. An outsourced model therefore transfers cost from build to specification and review, typically fifteen to twenty per cent of a senior person's time for the duration of the engagement, held by whoever owns the design. Confidence: Medium on that figure, which is drawn from comparable engagements and depends on how complete the specification is at handover.

Three risks weigh against the saving, and they are specific to this solution. Access to production data, or to the segregated identity store, from outside the EEA runs against the data residency position recorded in section 6, and brings transfer mechanisms and processor agreements that the in-house option does not; the mitigation is to withhold production access and work against synthetic data. The disclosure audit in phase 5 requires someone holding the whole anonymity model in mind, which is the work least suited

to a team engaged on volume. And intellectual property assignment runs through the vendor's contract chain rather than directly from the individual, which should be confirmed before engagement begins. Time zone overlap and turnover on the vendor side apply as they would on any engagement.

Elapsed time is the part of the comparison most often left out. The rate applies from the day development starts, and reaching that day takes months. Vendor identification and shortlisting come first, then scoping conversations and reference checks, then the commercial and legal work covering the contract, the data processing agreement, intellectual property assignment and any security review. Team assignment follows, and vendor availability rarely aligns with the buyer's timetable. Onboarding and ramp come last, and ramp is longer here than on a conventional build because the difficulty sits in judgement. Two to four months before the first line of code is a realistic planning figure, and that assumes no false start.

A nineteen to twenty-two week build therefore becomes a six to nine month elapsed programme, which is the figure that matters to a founder managing runway. The specification must also be complete enough to hand over before selection concludes, and that work does not presently exist; it falls on whoever holds the design, in advance, and is not costed in the ranges above. The in-house option can begin within weeks of a decision. Where time is the binding constraint the rate comparison does not decide the question. Confidence: Medium. The range reflects observed vendor selection cycles and was not modelled for this project.

A workable structure is to retain the phase 1 investigations, the technical design and the phase 5 disclosure audit in-house, and to outsource the implementation volume in phases 2 to 4 once the design is settled. The lead time is not fully additive in this structure. Vendor identification, scoping and contracting run alongside the phase 1 investigations and the in-house design work at the start of phase 2, which precede outsourced implementation in any case. That recovers part of the delay but not all of it. The structure therefore trades elapsed time for a lower rate. Section 1 records both budget and schedule as binding on this project, so which to favour is the company's decision. The observation this review can make is that the saving is uncertain, for the reasons above, while the delay is not.

Variance against the earlier estimate. An internal draft estimated 24 to 28 weeks. The lower figure in this report reflects the scope reductions in section 3.7 and the implementation of verification as a human workflow. Reinstating those items restores the five to six week difference.

Recommended review point

End of phase 3, at ten to twelve weeks of build, with about half the budget committed. At that stage verified experts are published, search operates under enforced anonymity rules, and enquiries are being received, with introductions completed manually. The expenditure that is difficult to reverse falls within phase 4: held funds, payouts, dispute handling and the calendar. Suspension or redirection at the end of phase 3 carries limited cost; after phase 4 commences it does not. The plan should be structured so that pausing at this point is a supported outcome.

Outstanding decisions affecting delivery

Whether experts are charged for listing, or the platform retains a share of each booking. This must be resolved before phase 2, since it determines whether seller income data is collected at expert registration, and the phase 4 build depends on it. The two models are distinct systems rather than alternative prices. A listing fee requires recurring billing and stored payment credentials. A share of each booking requires funds held between booking and delivery, deferred payout, refunds, dispute handling and income reporting. Implementing one and converting to the other subsequently would repeat most of phase 4. This review takes no position on the commercial merits. It records that the decision determines the six to seven weeks of phase 4 as well as the data collected in phase 2.

Selection of the identity verification vendor, and the policy for individuals not covered by any register.

Required before phase 2 concludes. The vendor selection is reversible. The policy is not, as it sets the standard applied to every expert verified to that point.

Whether permanent anonymity is reintroduced after the initial release. This review recommends deferring it from the initial release (section 3.7), and the estimates assume that. The recommended data model already separates identity, so reintroducing the feature later would not require the data model to be reworked. Reintroduction requires a supportive legal opinion, which need not be obtained before the initial build.

Operating costs

At launch scale, €85 to €395 per month. That is the sum of the fixed monthly items in section 3.4: hosting and database, search, transactional email, real-time messaging, file storage, and the identity verification vendor's monthly minimum where one applies, with monitoring and authentication on free tiers at this volume. Per-transaction and per-check charges, and the payment provider's per-connected-account and payout fees, are additional and scale with use. The figure is a list-price figure. In practice the first year is frequently covered by provider credits: a new cloud account presently carries a small starting credit, and the major providers operate startup programmes granting substantially more to eligible companies. Those credits defer the cost rather than remove it, and they expire at roughly the point the twelve-month figure below applies. A plan that models the first year on credits and does not carry the figure beyond them understates the running cost at precisely the moment revenue is expected to be carrying it. At twelve months, estimated between the client's six- and twenty-four-month targets, €800 to €2,000. The cost steps to monitor are the messaging service, which prices on concurrent connections, and the verification vendor, whose per-check charges overtake its monthly minimum as registrations grow. Neither warrants a different selection at this stage.

Exclusions

Design beyond the material already produced, content and copy, legal work, verification reviewer salary, the external security review recommended in section 3.10, marketing and any support function. The legal provision is not immaterial on this solution, for the reasons set out in section 6.

5. Risk register

Risk	Likelihood	Impact	Mitigation	Resolution point
Card authorisations expire before the session date	High	High	Capture at booking with deferred payout. Confirm during the first investigation.	End of week 1
Verification cost per expert is not established	High	High	Tier verification as set out in section 3.5 so the registry check is selective. Measure the reviewer effort for a manual check, record the external turnaround separately, and total the components before the fee is set. The proportion of experts needing manual corroboration must also be	Before the verification fee is set

Risk	Likelihood	Impact	Mitigation	Resolution point
			established, as appendix C records. Pricing itself is outside this review.	
A single verified badge implies more than was checked	Medium	High	Display identity, credential and experience separately, each stating its evidence. Decide the wording before the profile is built.	Phase 2
Identity disclosed through invitations, receipts or filters	Medium	High	Disclosure audit as a defined phase 5 task, with a checklist per surface.	Phase 5
Search returns insufficient results at launch	High	High	Build discovery so that it can be opened by field rather than all at once, and gate opening on a density threshold the company sets. The recruitment strategy that achieves that density is a commercial matter outside this review.	Pre-launch
Stale availability data offers unavailable slots	Medium	Medium	Defer availability filtering from the initial release. Definitive check at booking.	Phase 4
Specification and demonstration have diverged	High	Medium	Half a day of realignment before build commences. See also the limitation recorded in section 8.	Immediate
Funds cannot be held under a provider authorisation	Medium	High	Confirm in writing during phase 1. If the platform requires its own authorisation the payment design and the timeline both change.	End of week 1
No defined recovery objective for the identity store	Medium	High	Set a recovery point and time objective in phase 2, and test a restore that includes the segregated store during the phase 5 hardening work.	Phase 5
Data residency undermined by subprocessor or support access	Medium	High	Confirm EEA residency at subprocessor level, including support access, for each	Phase 1 for the payment provider, phase 2 for the rest

Risk	Likelihood	Impact	Mitigation	Resolution point
			vendor before contracting.	
Vendor selection delays the start beyond the runway	Medium	High	If an outsourced model is under consideration, begin selection in parallel with phase 1. Do not plan on development starting inside two months.	Before phase 2
Invoicing requirements conflict with anonymity	Medium	High	Settle the invoicing route during phase 3, before phase 4 commits the data model. Each route produces a different one.	Phase 3
Assessment based on the wrong artefact	Low	High	Confirm with the client which of the specification and the demonstration is authoritative, before build commences.	Immediate
Key-person concentration	Medium	Medium	The privacy and data model role is deliberately not rotated, as section 3.10 sets out, so the exposure is accepted rather than removed. Mitigate it by keeping the anonymity model documented to a standard a second engineer could pick up, and by walking a second engineer through the disclosure design at the end of each phase.	Ongoing

6. Regulatory considerations

The areas below are identified and characterised for planning purposes; they are not advised upon. Where an area is marked as applicable, professional advice should be obtained before the phase it affects. The conflict between erasure and financial retention shapes the data model and should be settled before phase 2. The holding of client funds is addressed first through the provider confirmation in phase 1. Where that shows the provider's model does not cover the activity, advice should be instructed immediately, because the answer can change the booking design built in phase 4; otherwise advice should confirm the position before phase 4. The invoicing route gates phase 4 and can be resolved during phases 2 and 3 provided it is instructed early enough. Regulatory observations are included to identify technical dependencies affecting delivery and architecture decisions.

Area	Applicable	Implication for the build
GDPR	Yes	Lawful basis per processing purpose, processor agreements with each vendor, subject access and erasure, and a record of processing activities.
Data residency	By election, not by law	GDPR does not require personal data to remain in the EEA. EEA hosting avoids the transfer question rather than satisfying a requirement. It constrains vendor selection. See below.
Special category data	Partly	The platform does not set out to hold health data, but experts in clinical fields may introduce it through messages or profiles. Policy should be established before launch.
Erasure against financial retention	Yes	An erasure request conflicts with transaction record retention obligations. Pseudonymisation of booking and payment records in place of deletion is the usual approach and should be confirmed with the adviser, with the policy documented.
Holding client funds	Likely	Capture at booking with deferred payout means funds are held for a period. A provider connected-account model ordinarily retains this activity under the provider's authorisation rather than the platform's. Confirmation is required.
Online platform obligations	Likely	An intermediary connecting buyers and sellers attracts notice-and-action duties and trader traceability requirements. Applicable thresholds should be established.
VAT and invoicing	Yes	Place-of-supply treatment, VAT number capture and validation, and whether the platform invoices as principal, issues invoices in the expert's name, or discloses the expert at invoicing. Only the first preserves anonymity at invoicing, and it changes the platform's VAT position. See section 3.2.
Seller income reporting	Likely	Platforms facilitating payment to sellers are generally required to report seller earnings to the revenue authority. This

Area	Applicable	Implication for the build
		determines the data collected at expert registration, depends on the revenue model decision recorded in section 4, and must be designed in early.
Employment and tax status	Unclear	Payment of experts through the platform raises questions as to their status. Advice should be obtained before the payout design is fixed.
Message scanning	Yes	Scanning of private messages for off-platform contact must be disclosed in the terms of service and privacy notice, with a stated lawful basis.
Advertising consent	Yes, if implemented	Targeted advertising on a free tier requires a consent management mechanism designed in from the outset.

Status of this section: the above identifies the areas that apply and characterises each one sufficiently for it to be planned around. It is not legal advice, it has not been settled with an adviser, and the position in each should be confirmed before it is relied upon. The legal provision for this solution should be set higher than is typical at this stage.

On data residency

This section assumes the operating company is established in Ireland and that users on both sides are within the EEA. If either changes the position below changes and should be re-examined.

The EU hosting decision is commonly misdescribed. The usual assumption is that personal data must be stored in the EU. The rules in fact govern the basis on which personal data is transferred outside the EEA rather than prohibiting transfer, so EEA hosting removes the transfer question rather than meeting a storage requirement. For this build the distinction has two practical effects. EEA hosting has been elected and the hosting, search and storage selections in section 3.4 assume it, so it constrains vendor selection for the initial release and should be treated as fixed. It also changes the cost of revisiting the decision later: a move outside the EEA would need a transfer mechanism rather than being ruled out. Whether a mechanism would be available on acceptable terms was not examined.

A related premise should be tested with an adviser, since it bears on any later move away from EEA hosting. It is sometimes assumed that data which never resides in the EU cannot be transferred out of it, so that hosting outside the EEA from the outset avoids the transfer rules. Published European Data Protection Board guidance indicates otherwise: the test turns on whether a controller subject to GDPR makes data available to a party in a third country, and not on where the data previously sat. Two further points for that adviser: transfers to the United States currently rest on the EU-US Data Privacy Framework, which is in force and under appeal, so standard contractual clauses should be held as a fallback; and EEA hosting is achieved only where subprocessors and vendor support access are EEA-resident, which is the most common way the position is undermined without anyone noticing. A full analysis of these areas is a due diligence exercise and was not undertaken within this engagement.

Priority: the holding of client funds should be resolved first, through the provider confirmation described at the start of this section. If the platform requires its own authorisation rather than operating under a

provider's, both the payment design and the delivery timeline change. The other areas follow the timing given there.

7. Summary for investor materials

The platform is a conventional web application constructed from proven components, with no dependency on unproven technology. A technical feasibility review found it feasible with constraints. It is achievable against the technology, subject to defined changes of scope and sequence, and its delivery depends on securing the budget, the elapsed time and two senior engineers. The review identified three design areas requiring deliberate handling: holding funds between booking and delivery, selective anonymity, and the verification of experts. It set out a phased delivery plan with a review point at ten to twelve weeks, at which verified experts and operational discovery are live and the model can be tested with users before the payment system is built. Measured in weeks of build, excluding any hiring or vendor lead time, the review estimated an early release at ten to twelve weeks, with introductions handled manually, and a full soft launch at nineteen to twenty-two weeks, with two engineers after a one-week investigation phase, at a build cost of €95,000 to €135,000. That figure covers development only. It excludes product and visual design, content and copy, legal work, the verification reviewer's salary, an external security review before launch, marketing and any support function, so the cost to launch is materially higher.

8. Limitations and exclusions

Commercial matters, entirely outside scope

- Market size, demand, and whether the demand side will pay for expert time. None of these was assessed.
- Pricing of any description: verification fee, commission rate, subscription tiers, corporate seats.
- Competitive position. A competitor analysis was provided and was read for context only. It has not been verified.
- The revenue model and financial projections.
- Route to market, and the accuracy of the outreach contact list.
- Company name, brand and domain availability.

Technical work not undertaken

- No source code exists, so no code audit was performed.
- No security audit or penetration testing, and no assessment of the non-functional requirements. Availability, resilience, security, maintainability and observability were not rated. Scale was considered only to the extent of establishing that the projected volumes require no unusual infrastructure, and recoverability only to the extent of noting the absence of a defined objective for the identity store; neither was rated as a non-functional requirement. A technical due diligence engagement covers these; a feasibility review does not.
- The regulatory areas in section 6 are identified and characterised for planning rather than analysed. Establishing the position in any of them is a due diligence exercise requiring legal input.
- No load testing or benchmarking.
- No accounts were created with, and no testing performed of, any vendor named in section 3.4.
- No third-party contract, licence or data-sharing agreement was reviewed.
- The capability of the team to execute was not assessed.

- The user interface specification was treated as superseded by the demonstration. Section 9 sets out why, and what changes if that choice is wrong.

Review period: five working days. The findings reflect what could be established within that period. Where the period was a constraint, the confidence rating records it.

9. Observations outside the assessment scope

The following fall outside the matters this review was engaged to address. They are observations rather than findings and do not carry the weight of the assessment above. They are recorded for completeness.

- **The specification and the demonstration describe different products.** The user interface specification documents an earlier generated screen set, with different expert names, prices in US dollars and layouts that no longer correspond to the current screens. The demonstration is a later rebuild of ten screens against a shared stylesheet, and is itself internally inconsistent, carrying one expert name on the dashboard and a different one on the profile. This assessment used the demonstration together with the product requirements document, and treated the user interface specification as superseded. That choice should be confirmed: if the specification is the intended reference, parts of sections 2 and 3 were assessed against the wrong artefact. Realignment requires about half a day.
- **A validation plan exists and has not been initiated.** It comprises interview scripts, defined thresholds and a concierge stage. Execution is not technical work and this review takes no position on its likely findings. It can be run in parallel with phase 2; its cost was not assessed here. Ownership is unassigned.

Appendix

A. Assumptions

The estimates assume the operating company is established in Ireland and that users on both sides are within the EEA at launch; phase 1 carried out by one senior person while the rest of the team is hired, and phases 2 to 5 by two full-time engineers at Irish contract rates; that the platform retains a share of each transaction rather than charging a listing fee; that the specification is substantially stable from the commencement of phase 2, with the level of change usual on a first product allowed for within the upper bound; that availability-filtered search and permanent anonymity remain deferred from the initial release; that the demonstration rather than the user interface specification represents the intended product; that verification is tiered as set out in section 3.5, with the formal credential check applied selectively; that the selected payment provider's connected-account model supports holding funds between booking and delivery for the entity type to be registered; and that verification retains a human reviewer. Should any of these prove incorrect, this is the list against which the estimates should be revisited.

B. Material reviewed

Two working sessions with the founder. A product requirements document covering user types, profiles, search, messaging, booking, payments, trust and safety, architecture and roadmap. Two versions of a user interface specification. A ten-screen click-through demonstration. A market research document. A validation plan. A competitor analysis and an outreach contact list, read for context only. Published pricing pages and published terms for the vendors named in section 3.4, and for HEDD, QQI and the shared digital credential platform, reviewed during the engagement. Published guidance consulted: card scheme and provider documentation on authorisation periods, regulatory guidance on Strong Customer

Authentication, European Data Protection Board guidance on international transfers, 2026 Irish contract rate guidance, published European ranges for senior advisory and fractional-CTO work, and published agency rate ranges. No source code, contracts or vendor discussions. The user interface specification was read but treated as superseded by the demonstration, for the reasons in section 9.

C. Unverified inputs

The following were relied upon during the assessment but not independently verified. They are recorded so that their status is clear and so that confirmation can be prioritised.

Input relied upon	Source	Required to confirm
The authorisation window, standard and extended	Scheme documentation and provider guidance. Standard and extended periods, eligibility and any per-transaction fee vary by scheme, acquirer and merchant category.	Written confirmation from the payment provider ultimately selected, covering whether extended authorisation is available to this entity type and on what terms.
The two to four week booking lead time for specialist sessions	Stated in the client's requirements and used as the basis of the payment timing analysis. No market data was examined.	Observed booking behaviour once live, or evidence from a comparable marketplace.
The Strong Customer Authentication treatment of charges made without the cardholder present	Published regulatory guidance, read for the purpose of comparing payment designs.	Confirmation from the payment provider selected, for the integration chosen.
All vendor pricing in section 3.4	Published pricing pages reviewed during the engagement.	Written quotations at projected volumes.
Per-institution verification fees and turnaround for Irish universities	QQI and HEDD published terms were checked, and the shared credential platform's published institution list was read; individual university registry services were not contacted and the platform's volume terms were not obtained.	A written service level from each institution in the initial launch set, and commercial terms from the credential platform.
Infrastructure cost at twelve months	Estimate derived from comparable systems, not a model of projected traffic.	A capacity estimate once message and search volumes are established.
The six and twenty-four month scale targets	The client's own planning documents. No supporting evidence was examined.	A bottom-up model, or observed traffic once live.
That most expert categories are not covered by any register	Assessment of the stated category list against known registers.	Review of the list with a specialist recruiter in each field.

D. Glossary

Authorisation and capture. Two distinct stages of a card payment. Authorisation confirms funds are available and reserves them. Capture collects them. The interval between the two is limited, which is the constraint described in section 3.2.

Connected account. An arrangement under which a payment provider maintains an account on behalf of each seller, allowing a platform to collect and disburse funds without handling them directly.

Denormalised index. A secondary copy of data, restructured for rapid search. Fast to query, and subject to drift from the source record.

Initial release. The product as delivered at soft launch, at the end of phase 5. Distinct from the early release at the end of phase 3, at which introductions are still handled manually.

Technical investigation. A short, time-boxed investigation addressing a defined question or small set of questions, with no expectation that the resulting code is retained.

State machine. A design in which an entity may occupy only one of a defined set of states, with permitted transitions between them specified. Applied here to prevent a booking reaching an inconsistent part-paid, part-cancelled condition.