Pre-Hire Methodology Brief

A staged pre-employment screening service: what it measures, how the composite is scored, what the screening report is for, and the one rule that governs every part of it - VIFM recommends, the client decides.

Version 1.3 · 17 November 2026

1. Why this exists

Whenever an assessment informs a hiring decision, three questions arise: *What does the instrument actually measure? How is a single number derived from a candidate's answers? And what weight should the number be allowed to carry?* This brief answers those questions for Pre-Hire, the pre-employment screening service through which VIFM offers the platform's existing assessment instruments to client organizations as a structured, defensible shortlisting layer.

Pre-Hire produces a screening signal that supports a human decision. It is deliberately not a hiring decision, not an automatic rejection, and not a standalone verdict on a candidate. The single most important design choice in the entire module is this: **the scoring logic has no rejection outcome.** A candidate is never declined by software. VIFM is the assessor and delivers a report; the engaged client makes the hiring decision. Every other design choice documented below follows from that positioning.

2. What Pre-Hire measures

An orchestration, not a new instrument

Pre-Hire does not introduce a new test. It is an orchestration layer over instruments that already exist on the platform, assembled into a client-specific, role-specific screening sequence. Each stage in a requisition is scored by the existing instrument's own method rather than by a separate re-implementation. The three interactive screening stages are:

- **Competency Quiz** - a competency and knowledge screen produced by the VIFM quiz generator and delivered in English and Arabic. Each selected competency is covered by three questions (up to ten competencies). Questions are drawn first from VIFM's expert-approved question bank; where the bank does not yet hold enough approved questions for a competency, the remainder are generated by AI. While the approved bank does not yet cover every competency, the screening report is marked as preliminary. Answers are graded on the platform, and the answer key is never sent to the candidate's device. The normalized stage score is the percentage of questions answered correctly, rounded to the nearest whole number.
- **English (Fluent)** - the four-skill English placement (reading, listening, writing, speaking). A requisition may administer a subset of the four skills, provided that reading or listening is included; the report then labels the result as a partial placement. Reading and listening are auto-scored; writing and speaking are scored against a CEFR rubric, optionally blended with pronunciation assessment. The overall CEFR level is converted to the 0-100 scale in equal steps of 20: A1 = 0, A2 = 20, B1 = 40, B2 = 60, C1 = 80, C2 = 100.
- **AI Interview (CBI)** - a structured competency-based interview conducted by the AI interviewer. For new requisitions, the interview covers the requisition's highest-priority competency and its approved behavioral indicators: an opening question and up to three follow-up probes (situation, task, action, result), so at most four candidate answers. Requisitions already in progress keep the interview they started with, which covers one role-level behavioral competency named after the role. The transcript is rated by AI on a single 1-5 behaviorally anchored rating scale (BARS), converted to the 0-100 scale in equal steps of 25: 1 = 0, 2 = 25, 3 = 50, 4 = 75, 5 = 100. A VIFM reviewer can review and certify the result.

A requisition's stage plan is an ordered list of stages, each carrying a stage type, a weight, a cut-score, and a Required setting. Each stage can be marked Required when the requisition is set up or edited. A candidate below a required stage's cut-score is never recommended to advance (see §3). Which stages start as Required depends on how the requisition is created:

- **Set up by VIFM for a client:** every stage starts with Required switched off, and VIFM switches it on for the stages the client wants to treat as must-pass.
- **Set up by the client from its own portal (and in the demo):** unless VIFM has set a plan for that client, the standard plan applies, which weights the Competency Quiz 0.4 (cut-score 60), English (Fluent) 0.3 (cut-score 50) and the AI Interview 0.3 (cut-score 60), with the Competency Quiz and English stages marked Required.

In every case the stages, weights, cut-scores and Required settings are configured for each role.

Defining "good" per role

What counts as a strong screening signal is defined per role, not globally. Each requisition can be linked to a role profile and an explicit competency set, from which the quiz questions are drawn, and carries its own cut-scores and weights, so the quiz items, the cut-scores, and the weighting all reflect the specific opening. This per-role configuration is intended to make the composite relevant to the specific opening rather than a generic number. The cut-scores and weights are set by judgment; they have not yet been calibrated against job performance data.

3. How it is scored

From stages to a composite

Each completed stage contributes a normalized 0-100 score. The composite is a weighted average of those stage scores. Only stages with a positive weight contribute. Weights are normalized against the total weight of the weighted stages, so they need not sum to 1: the weight applied to a stage is its weight divided by the total, and the composite is the sum of each weighted stage's score multiplied by its applied weight, limited to the 0-100 range and rounded to one decimal place.

The composite is **not produced until every weighted stage and every required stage has been scored.** A partial profile produces no headline number. This is deliberate: a weighted average over an incomplete set of stages would understate or overstate a candidate, so the module shows no composite at all until the picture is complete. A required stage that carries no weight must therefore still be completed, even though it does not change the number.

Per-stage pass and fail against the cut-score

Each stage is independently marked against its cut-score. A scored stage with no cut-score always passes; a scored stage with a cut-score passes when its normalized score is at or above the cut-score. A stage is shown as not yet evaluated until it has been scored. A stage marked as required that falls below its cut-score is recorded as a required-stage shortfall.

The advisory recommendation (and the absence of "reject")

From the composite and the required-stage check, the platform derives one of exactly four advisory bands - Advance, Review, Hold for review, and In progress - using fixed thresholds on the rounded composite (advance at 70, review at 50). The rules are applied in this order:

1. **In progress** - the candidate has not yet been scored on every weighted stage and every required stage, so no composite exists yet.
2. **A required stage is below its cut-score** - the signal is capped. It is **Review** when the composite is 50.0 or above, including when the composite is 70.0 or above, and **Hold for review** when the composite is below 50.0. It is never Advance, and the candidate is never dropped.
3. **Advance** - composite at or above 70.0 and every required stage at or above its cut-score.
4. **Review** - composite from 50.0 to 69.9.
5. **Hold for review** - composite below 50.0. "Hold" means "low signal, a person should review," never "rejected."

A stage that is not marked as required can fall below its cut-score without capping the signal; the report still shows it as below its cut-score.

There is no fifth band. A required stage below its cut-score caps the advisory signal at Review or Hold for review but never produces a decline. The composite, the advisory band, and the candidate's status are recalculated and saved after each stage is completed, and the shortlist is ordered from the highest

composite to the lowest, with candidates who do not yet have a composite placed last. The ranking orders candidates for a human reviewer's attention; it does not exclude anyone.

---

4. Deliverables

The per-candidate screening report

The primary deliverable is a per-candidate screening report, available in English or Arabic and produced as a PDF. VIFM is the assessor, not the client, so this is the document VIFM downloads or emails to the engaged client. It presents:

- The advisory composite and the recommendation band, color-coded for quick reading.
- A by-stage table: each stage's weight, normalized score, cut-score, and pass / below-cut / not-taken outcome.
- A "How this score is calculated" section that states the composite formula and lists all four advisory bands with the condition that triggers each one. Because the signal is derived automatically, the client is told how.
- An explicit disclaimer on the face of the report: this is an advisory screening signal, not a hiring decision; a qualified VIFM reviewer interprets it alongside other evidence; no decision is ever made automatically.

The report is also marked confidential, for VIFM and the engaged client only.

ATS export

A structured export (JSON or CSV) is available to VIFM administrators for passing results to a client's applicant-tracking system (ATS). The composite is recalculated at the time of export, so the export always reflects the current stage results, and the CSV is encoded so that Arabic and accented names display correctly in spreadsheet tools. The export deliberately contains no demographic data.

---

5. What Pre-Hire is, and what it is not

- **It is a screening signal that supports a human decision.** It combines several existing instruments into a transparent, role-weighted composite to help a reviewer prioritize attention. That is its purpose, and the full extent of its role.
- **It is not a hiring decision and not an automatic rejection.** The scoring logic has no rejection outcome. The advisory band caps how a screening conversation starts; it never closes one. A person makes every decision, and VIFM delivers a report - the client decides.
- **It does not re-implement the underlying instruments.** Each stage is scored by the existing instrument's own method. Pre-Hire orchestrates the stages and combines their results.
- **Fluent is a screening signal until it is calibrated and proctored.** By design, the English placement is treated as an indicative screen, not a hard gate, until calibration evidence (target quadratic weighted kappa, QWK, at or above 0.70 against human raters) and a proctoring decision are in place. In the default plans the English stage is required with a cut-score of 50, so a Fluent result below the cut-score caps the advisory band at Review (Hold for review when the composite is below 50); it never rejects a candidate on its own.
- **Predictive validity is not claimed before it is evidenced.** VIFM's development plan makes a completed validity and fairness study a condition of offering the service beyond its pilot; the service does not claim a prediction it has not yet demonstrated.

These limitations are not an afterthought. The "screening signal, not a hiring decision" line is printed on the face of every report, in English and in Arabic.

---

6. Defensibility: decoupled demographics, adverse-impact monitoring, and an immutable audit trail

The defensibility layer is what makes Pre-Hire suitable for use in a selection context, and it has three essential parts.

Voluntary, decoupled demographics

Self-identification is optional, appropriate to the GCC context (gender, age band, and national or expatriate status), and entirely decoupled from scoring. It is requested in its own voluntary step, immediately after consent and before the first assessment stage, so it is visibly separated from any assessment content or result, and it is never an input to any score. The candidate may skip the step, and "Prefer not to say" is an explicit option for every question. Individual demographic values never appear in the audit trail, the ATS export, or any client-facing surface; they exist only to support the aggregate fairness view, and they are never imputed.

Adverse-impact monitoring (the 4/5ths rule)

The fairness view applies the standard four-fifths screen used in employment selection. For each demographic dimension it computes the selection rate per group, takes the highest-selected group as

the reference, and flags any group whose impact ratio (its rate divided by the reference rate) falls below 0.8. The framing is explicit and built into the analysis:

- It is a **monitoring signal, not proof of discrimination.** A flag warrants reviewing the instrument's job-relatedness, not an automatic change.
- **Small samples are surfaced, not hidden.** A group smaller than five members is marked as a small sample, and a dimension is marked underpowered when the pool is below thirty (a practical minimum VIFM applies; the Uniform Guidelines caution that differences based on small numbers may not be meaningful, without setting a number), when any group is small, or when there are fewer than two groups. The fairness view shows the group sizes and the underpowered caveat rather than implying false precision.
- **Recorded decisions first, the screening recommendation otherwise.** The client or VIFM records the hiring decision for each candidate (Advanced or Not advanced, with an optional reason). Where decisions have been recorded, the fairness view uses them, so it shows whether hiring outcomes differ between groups. Where no decision has been recorded, it falls back to the advisory band, counting a candidate as selected when the band is Advance; it then shows whether the screening signal differs between groups, not whether hiring outcomes do. Non-disclosed demographics are excluded from group rates rather than estimated.

Immutable audit trail

The platform keeps an append-only record of significant actions: requisition created or updated, candidate added, invitation sent, consent given, stage completed, demographic step submitted, reviewer certification recorded, report shared, and export taken. Entries cannot be edited or deleted once they are written. The audit record never contains demographic values, so the trail can be made readable to the client without exposing personal self-identification data. Audit logging never blocks or interrupts an assessment: if an audit entry cannot be written, the assessment continues.

---

7. Access model, fairness, confidentiality and retention

Candidate access without an account

Candidates have no platform account. Each candidate receives a unique, private access link, and the platform identifies the candidate from that link on the server at every step of the flow. Administrative functions such as the ATS export are available only to signed-in VIFM administrators and are kept separate from the candidate flow, so they cannot expose candidate data.

Standards the service is held to

Pre-Hire's design and intended use are aligned with the following recognized standards for the proper use of assessment in selection. Alignment is VIFM's own mapping against the published standard or law; it is not a certification, accreditation, audit or legal opinion by the body concerned.

- **ITC Guidelines on Test Use** - on the responsibilities of those who select, administer and interpret assessments.
- **Standards for Educational and Psychological Testing** (AERA / APA / NCME, 2014) - for the discipline of matching claims to evidence and for the validity and fairness evidence the pilot is built to gather.
- **Uniform Guidelines on Employee Selection Procedures (1978)** - the four-fifths adverse-impact screen follows the method described in the Uniform Guidelines, and a person, not the software, makes every selection decision.

Proper use and confidentiality

The human-in-the-loop control is structural, not cosmetic. There is no rejection outcome in the scoring logic; the report states on its face that the signal is to be interpreted alongside other evidence before any decision is made; the ATS export carries no demographic data; and the fairness view is restricted to VIFM administrators and shows aggregate figures only. The screening report is a controlled deliverable that VIFM hands to the engaged client; per-candidate scores and answers are handled as confidential personal data throughout.

Retention

Pre-Hire data follows the platform's default retention rule - a maximum of two years unless a contract extends it - designed with reference to UAE Federal Decree-Law No. 45 of 2021, the Saudi PDPL, and GDPR where applicable, and to the platform's consent-before-collection commitment. Consent is captured as a gate at the start of the candidate flow, before any stage runs, and the consent event is recorded in the audit trail.

---

8. Bilingual availability

Pre-Hire is available in English and Arabic. The candidate interface and the Competency Quiz are offered in both languages; the English (Fluent) stage is, by design, an English-language placement; and the AI Interview is currently conducted and scored in English. The screening report is produced in English and in Gulf-appropriate Arabic, mirroring the same underlying data - the same composite, the same per-stage scores, the same band methodology - so a report reads identically in either language and a bilingual cohort is scored on one consistent basis. The CSV export is encoded so that Arabic names display correctly in spreadsheet tools. Arabic content is provided on a best-effort basis, pending human review before high-stakes use.

---

9. Contact

For methodology questions or research-collaboration inquiries, contact the VIFM consulting team: `contact@viftraining.com`.

Reliability and Validity Evidence

Two kinds of evidence: **(A)** established coefficients for the *methods* combined in the composite, from the published meta-analytic literature; and **(B)** VIFM's own reliability and defensibility procedures with their current status.

A. Established coefficients for the methods (published literature)

- **Structured / behavioral interview** (the CBI stage): published meta-analyses of structured employment interviews conducted and rated by trained human interviewers report corrected validities of about **.44** (McDaniel, Whetzel, Schmidt & Maurer, 1994) to **.51** (Schmidt & Hunter, 1998), against about .33 and .38 for unstructured interviews in the same two sources. The Pre-Hire AI Interview is a short, AI-conducted and AI-rated interview on one competency; these published figures describe the general method and have not been shown to apply to it.
- **Composite validity:** published research on combining selection methods (Schmidt & Hunter, 1998) reports that adding a second, different method to a general mental ability test raises validity. Pre-Hire's composite combines a competency quiz, an English placement and an AI-conducted interview; it does not include a general mental ability test or a work sample, so this finding does not transfer directly, and the validity of the Pre-Hire composite has not yet been established.
- The Fluent placement is described in its own methodology brief. The competency quiz and the AI interview do not yet have separate briefs; their evidence is limited to what is stated here.

B. VIFM's own reliability and defensibility - method, threshold, and current status

- **Reliability not yet estimated:** reliability of the composite has not yet been estimated. Planned: internal consistency of the quiz, agreement between AI and human ratings for the interview and for Fluent writing and speaking, and the stability of the composite on re-test.
- **No items of its own:** Pre-Hire is a **weighted composite** with per-stage cut-scores. It produces an advisory band (advance / review / hold) and is **never an automatic rejection - a human always decides**.
- **Adverse-impact monitoring:** the **4/5ths rule** (Uniform Guidelines on Employee Selection Procedures, 1978) is computed live per demographic dimension; an append-only audit trail records significant actions.
- **Criterion validity** of the composite is not claimed pending a local predictive study.
