HIHB

EN EDITION - 2026

STUDY

State of Hiring Briefings 2026

The workshop delta in 44 Pre-Recruiting mandates.

Michael von Hirschfeld, HireWorks GmbH
June 2026

Method: High-Impact Hiring Blueprint (HIHB), Pre-Recruiting.
Data basis: content analysis of 44 anonymised recruiting canvas workshops (B01 to B44).

Contents

STUDY		PAGE
01	Foreword - the core of the study in one sentence	3
02	Research question and contribution	4
03	Data basis	5
04	Methodology	7
FINDINGS AND VIGNETTES		
05	Findings	9
06	Three stakeholder vignettes	11
07	Before-and-after metrics	14
RECOMMENDATIONS AND APPENDIX		
08	Recommendations	16
09	Appendix	18

CHAPTER 01

Foreword – the core of the study in one sentence

Author: Michael von Hirschfeld, HireWorks GmbH

Method: High-Impact Hiring Blueprint (HIHB), Pre-Recruiting

Data basis: content analysis of 44 anonymised recruiting canvas workshops (B01 to B44)

As of: June 2026

This study is the first to measure what a moderated briefing workshop actually adds to a key hire.

It does so not through the usual self-report of an HR department, but through the workshop artefacts themselves: forty-four fully completed recruiting canvas boards from eight years of consulting work, extracted verbatim, anonymised and systematically evaluated.

The central finding is not intuitive. The workshop effect does not lie in the number of notes a stakeholder trio produces in two hours. It does not lie in a broader picture of the role, nor in a longer list of requirements. It lies in the methodological depth of a few load-bearing fields. Where stakeholders without moderation would fill description fields, the workshop adds measurement fields. Exactly there, and only there, arises what this study calls the **workshop delta**.

> **Not "how much did we discuss", but "at which points did we measure instead of describe".**

The observation has two consequences. First, it changes the question to ask about Pre-Recruiting workshops. Not "how much did we discuss", but "at which points did we measure instead of describe". Second, it changes the briefing format itself. Four template generations from our own archive show how the method has moved: away from collecting, towards calibration; away from the requirements lump, towards the must/teachable split. This study makes this development visible and discloses the data on which it rests.

CHAPTER 02

Research question and contribution

Primary question

What does a moderated Pre-Recruiting workshop achieve that two hours of unmoderated self-reflection by the stakeholders would not?

Secondary questions

Which fields of a recruiting brief carry weight reproducibly, and which remain structurally thin? At which points do reproducibly measurable statements arise, and where does the open question produce wish vocabulary? Which pain patterns repeat across industries and template generations? From what sample size do no new finding categories appear?

Contribution to the Pre-Recruiting discussion

The established recruiting literature measures outcomes after the hire: time to fill, quality of hire, failure rates, tenure. It treats the brief as a black box. This study shifts the focus to the phase before. It provides a before-and-after measure of briefing quality and makes the workshop effect quantifiable for the first time.

We show three dimensions in which the effect can be operationalised: measurable success criteria per horizon, prioritised key activities with a hundred-percent distribution, and a must/teachable split of the requirements. Pre-Recruiting is thus no longer the thesis "start earlier" but the thesis "start in a methodically measurable way".

Scope

We do not rely on self-report or on outcome data. We read workshop artefacts. This has a limitation: we can argue the effect of the workshop on the later quality of the hire only indirectly, via the mechanism, not via the outcome. But it also has a strength: the data exists independently of the participants' memory and can be recounted reproducibly. External recruiting studies (SHRM Recruiting Benchmarking, Leadership IQ "Why New Hires Fail") appear in this study only in the recommendations chapter, as triangulation against outcome findings, not as a data basis.

CHAPTER 03

Data basis

The corpus comprises forty-four recruiting canvas workshops, documented internally as B01 to B44. The boards come from the complete HireWorks workshop archive as of 1 June 2026. This is not a sample and not a selection. It is the entire archive. Cherry-picking is therefore ruled out; individual boards cannot be swapped to support a thesis.

Every board was extracted verbatim. The first four boards (B01 to B04) were read and coded directly. The remaining forty were extracted by eight parallel subagents under strict verbatim and anti-fabrication discipline (see the methodology chapter). Real names exist only in the internal master. Every statement in this study uses the B-ID plus industry cluster, never company names, never personal names.

Four template generations

The corpus contains four template generations side by side, which documents how the workshop method developed over eight years.

- **Template A** is the HireWorks recruiting canvas with fields 1 to 9 (Current Situation, Structure, Key Activities, Experience, Culture Add, Core Questions, Success Factors, Compensation and Perspective, Recruiting and Qualification Process). It is the largest subset with twenty-eight homogeneous boards, plus two A-like special cases with a different field assignment.
- **Template B** is the older ThinkHeads recruiting canvas, recognisable by its green wordmark. Eight boards, mostly from the early phase.
- **Template C** is the "Hiring Requirements Engineering" template, which introduced the must/teachable split for requirements for the first time. Four boards.
- **Template D** is the High-Impact Hiring Blueprint, the most recent generation. It integrates persona, why, impact scoring and a must/nice split. Two boards.

Industry clusters

The primary assignment gives the following picture: digital and tech (product, IT, data, engineering) about fourteen boards. Marketing, media, sales and ad tech about eleven. Executive and department heads about six. Operations about four. Finance three. HR and recruiting two. The remainder is spread across mixed roles that cannot be clearly assigned.

Period

The boards date from 2018 to 2026. Individual boards with a COVID context come from 2020 and 2021. The template evolution is visible in the corpus: older boards in template B, the middle phase in template A, the most recent boards in C and D.

Limitations

Three caveats apply. First, the corpus leans towards digital and tech; classic industrial key roles are underrepresented, and the findings are calibrated to knowledge-intensive leadership and specialist roles. Second, field names and branding vary between template generations; we make longitudinal statements across all

forty-four boards only where the bucket logic holds. Third, the fill-level matrix captures the fields Current Situation, Key Activities, Success Factors and Requirements (CF1 to CF4) cleanly for only about thirty-one boards, because thirteen boards use a different note format. The findings on the less filled fields (CF5, CF6 and the non-canonical fields NC1 to NC3) rest on forty to forty-two boards and are robust.

CHAPTER 04

Methodology

The methodology is the heart of this study. It has three parts: the collection of the data, the anticipation method for the "before", and the quality assurance that supports the findings.

4.1 Collection

All forty-four boards were extracted verbatim. The subagents worked under a strict anti-fabrication protocol: no conclusions from guessed content, every dense or hard-to-read note marked as [illegible] or [partial], no interpretation during the extraction phase. The notes of each board were then mapped to a six-part common-field logic (CF1 to CF6) plus three non-canonical fields (NC1 to NC3). The fill-level matrix counts the number of notes per bucket per board. Qualitative coding of pain types, requirement archetypes and client language patterns took place in a second wave against the coded notes.

4.2 Anticipation method: the delta argument

The core of the study is the workshop delta. It is the difference between what stakeholders produced AFTER two hours of moderated workshop and what they would have produced WITHOUT a workshop in two hours of self-reflection. The latter cannot be observed directly - we have no control-group experiment in which the same stakeholders reflected alone for two hours. We anticipate the "before" in a methodically sound and transparent way. Three building blocks combine, each shown separately.

BUILDING BLOCK (A): INFERENCE FROM FILL LEVEL AND FIELD QUALITY

Fields with a high fill rate and little methodological depth would have come about in comparable form without a workshop. The reasoning is simple: what stakeholders produce in two hours of workshop, they can also produce in two hours alone - provided the field only asks for description. Specifically, this concerns two fields. **Current Situation** (CF1) has a median of thirteen notes per board in our corpus and is never empty. The field describes an existing situation: predecessor role, state of the team, reorganisation, interim solution. This information is in the stakeholders' heads; the workshop brings it out but adds nothing methodologically. The same applies to **Key Activities** (CF2): median eighteen notes, never empty. Stakeholders can list the tasks of a role without anyone asking. These fields are well filled, but they are description fields. They are not a workshop effect.

Fields with methodological depth, on the other hand, are the workshop delta. Four layers belong here and can be identified reproducibly in the data. First, the **must/teachable split** in the requirements, which appears systematically only in the more recent templates C and D. Second, **measurable success factors per horizon** instead of mere time markers - this step requires a moderation move that replaces wish vocabulary. Third, **critical-few reasoning and weighting** of the key activities, i.e. the hundred-percent distribution. Fourth, the **triad consensus markers** between CEO, HR and department head, where diverging views first become visible and are then negotiated. In our data, this methodological layer arises only through moderation; it is Pre-Recruiting in the narrow sense.

BUILDING BLOCK (B): STAKEHOLDER VIGNETTES

Three typical pre-workshop views can be reconstructed from the boards. CEO and management arrive with a strategic charge, vision vocabulary and a task list that pulls the role towards strategy. HR and talent arrive with a requirements lump without a must/teachable separation, years of experience as a proxy for suitability.

Department heads arrive with operational detail in the tasks, less connection to strategy, and an often very clear persona that is rarely documented. For each vignette it is possible to reconstruct plausibly what the role would have produced in two hours without moderation - and what the workshop added. The point of convergence where the three views meet is where the workshop delta arises. The three vignettes are written out in chapter 6.

BUILDING BLOCK (C): BEFORE-AND-AFTER METRICS

Three dimensions can be quantified. First, the share of measurable success criteria per horizon (three, six, twelve months) versus mere time markers. Second, the share of prioritised key activities with a hundred-percent distribution versus an unweighted collection. Third, the share of requirements with a must/teachable split versus the requirements lump. For each dimension, the pre-workshop anticipation (derived from building block a) is set against the post-workshop actual (measured from the forty-four boards). The quantification is in chapter 7.

A fourth dimension - the share of triad consensus markers - would be methodologically interesting but needs an additional coding pass and is not quantified in this version of the study. It is shown qualitatively in the vignettes.

4.3 Quality assurance

Four measures support the findings.

ANTI-FABRICATION PROTOCOL

The subagents extracted verbatim. Dense or hard-to-read sticky notes were marked as [illegible] or [partial], never guessed. This discipline is not trivial: in an early test, a subagent produced plausible-sounding note content that was not in the original. This led to a tightening of the protocol and an obligation to mark every unconfirmed passage.

B25 QUALITY CHECK

Sample B25 (template C, industry cluster finance) was checked field by field against the original image. The extraction matched exactly, including the colour coding of the sticky notes. The high internal coherence of all eight subagent groups confirms that the protocol scales.

SATURATION CHECK

The dominant pain and requirement categories repeat after an estimated twenty to twenty-five boards. New categories hardly appear after that. This is inference, not a formal code count per note. But the implication is robust: pattern recognition in Pre-Recruiting scales with moderately sized corpora; a hundred-board study would be contrived.

ANONYMISATION

Real names remain in the internal master. Every public statement uses the B-ID plus industry cluster. The stakeholder vignettes in chapter 6 are generic archetypes, not individual cases - that is a methodological decision, not a narrative weakness. We wanted to rule out re-identification structurally, even if the individual vignette loses some vividness as a result.

CHAPTER 05

Findings

The findings summarise the main observations of the content analysis. They are not exhaustive, but condensed to the points that shape the structure of the Pre-Recruiting discussion.

5.1 Four template generations - the method has visibly evolved

The corpus is template archaeology. The template distribution (A thirty, B eight, C four, D two) is not random; it reflects how HireWorks developed its own method over eight years. The development from ThinkHeads via the HireWorks recruiting canvas to Hiring Requirements Engineering and finally to the High-Impact Hiring Blueprint closes exactly the gaps that become visible in the older templates. ThinkHeads collects; HireWorks orders; Hiring Requirements Engineering calibrates; the Blueprint personalises and weights.

Three steps can be clearly identified methodologically. The first step is the introduction of the **must/teachable split** in the requirements, visible in template C. Before this step, all requirements clumped together in a single field; afterwards, the workshop separates "must be there on day one" from "can be built up in the first year". The second step is the introduction of **persona and why** in template D, i.e. the shift from "who do we want who meets this list" to "who is this person, in what life situation, with what drive". The third step is **impact scoring** and the distribution logic for key activities, anchored tentatively in C and systematically in D.

The template evolution is not just a historical detail. It is the chain of evidence that the method has evolved out of practice. Each new generation responds to a weakness that became visible in the previous one.

5.2 Three fields carry the weight, two are effectively dead

The fill-level matrix shows a clear split. Three fields carry weight reproducibly. **Current Situation** is empty in none of the boards captured; the median is thirteen notes, with individual peaks of thirty-five (B36) and twenty-five (B07). **Key Activities and Tasks** is also never empty; the median is eighteen notes. **Requirements** is always filled; the median is seven notes, already split into must/teachable in templates C and D.

Two fields are effectively dead. **Recruiting Process** is completely empty in twelve of forty-two boards captured - almost a third of the corpus. **Core Questions** or Qualification is empty in ten of forty boards - a quarter. Both fields fall below the fill-level threshold at which a field can be considered load-bearing. **Structure** lies in between: the median is three notes, often just an org chart, empty in a few boards.

The implication is methodologically direct. The dead fields are not a question of industry or role; they are spread across all clusters. It is a problem of method and moderation. Either the fields are redefined in a coming template generation (a different set of questions, a different place in the sequence), or they are dropped. The data does not support the middle way of "we hope the stakeholders will fill it in".

5.3 Success factors as a paper tiger

The **Success Factors** field has a median of eight notes per board. That suggests a quality that is not there. In terms of content, most boards contain either the bare time frame "3 months / 6 months / 12 months" without measurable content, or wish vocabulary: "good collaboration", "meet internal deadlines", "integrated into the team". These notes are not wrong. They are just not observable. A statement such as "good collab-

oration with the management team" cannot be checked after six months, because it names no measure against which success could be assessed.

Three boards illustrate the pattern. B03 fills the field only with the three horizon markers, with no content below them. B25 fills it with the wish phrases quoted above. B29 even leaves the twelve-month marker empty - the time frame itself remains incomplete. In contrast, B04 shows that the field can carry weight: here we find statements such as "get the current version running" and "design a completely new app" - pictures of success that can be observed after six months. It is not the field, it is the moderation.

The methodological lever is clear: enforce one observable measure per horizon. The question "How will the team recognise success after six months - in an observable measure?" delivers exactly that in the boards with high discipline (B04 as the positive example). Phrased openly, it delivers wish vocabulary.

5.4 Key activities tip over into overload

Without a stop rule, the **Key Activities** field becomes a wish list. The peaks in the corpus are striking: B28 with forty notes, B37 with twenty-nine, B08 with twenty-eight, B36 with twenty-four, B07 with twenty-one. With forty notes per role, the field can no longer serve its selection function. It becomes an inventory of everything the role holder "should also be able to do".

The diagnosis is not that the notes are wrong. The diagnosis is that they are unweighted. Without a stop rule and without a hundred-percent distribution, the field produces a false completeness that becomes expensive in the selection process: in the end, assessors look for candidates who meet the entire list and fail - or they ignore the list and decide intuitively. Both lead to a worse selection than a focused list of eight to ten critical-few activities.

Positive exceptions can be identified in the corpus. B08 and B28 show percentage weighting, i.e. an attempt at distribution. That is exactly the mechanism the more recent templates anchor systematically. The methodological lever: stop at around ten notes, then group, then distribute a hundred percent. Where the distribution grinds, the mismatch between the views becomes visible.

5.5 Pain typology

Across industries and template generations, four pre-workshop pain types repeat, visible in the **Current Situation**. First: the predecessor has left and a successor is to take over. Second: an interim solution is in place and the stopgap is to be replaced. Third: reorganisation or restructuring, the role has been newly created or reshaped. Fourth: scale-up and growth, the role does not yet exist in this form. Boards illustrate this clearly: B31 documents an active interim solution; B05 records the creation of the role in a reorganisation; B01 shows the growth logic of a scale-up. A fifth sub-category, the COVID context with short-time work and hiring freezes (visible in B31), appears only historically in the boards from 2020 and 2021.

The typology is not exhaustive, but it works as a conversation framework. In the workshop, it helps structure the open question about the starting point without forcing answers.

5.6 Saturation point at around board twenty to twenty-five

After an estimated twenty to twenty-five boards, the dominant pain, requirement and language categories saturate. New categories hardly appear after that. This is inference, not a formal code count per note, but the implication is methodologically important: a hundred-board study would be contrived. A corpus of forty to sixty boards is the right size for pattern recognition in the Pre-Recruiting phase.

CHAPTER 06

Three stakeholder vignettes

The vignettes show the convergence mechanics of the workshop from three perspectives. They are generic archetypes, not individual cases. Placing them in an industry cluster makes them concrete without making them re-identifiable. Each vignette shows what the person would have produced alone before the workshop, what the workshop added, and which concrete element made the difference.

6.1 Vignette A: CEO of a mid-sized tech company

Profile. CEO and co-founder of a B2B tech company with eighty employees, industry cluster digital/tech. After three funding rounds, the company is in its growth phase, and a new head of growth is to be hired. The CEO drafted the briefing document himself and passed it on to the HR department.

Before the workshop. The CEO arrives with strategic vocabulary. He describes the role as "the growth engine for the next three years". He has a task list of about twelve points in mind, roughly divided into strategy (sixty percent) and operations (forty percent). Asked about success criteria, he answers with the three-six-twelve-month markers and wish vocabulary: "trust built in the first quarter, first results in the second, growth team established after a year". He has not worked out the list of requirements; "HR should do that".

What he would have produced alone. An extensive current situation with vision vocabulary and a market description. A task list that quickly exceeds thirty points, because every further thought adds another type of task. Success criteria as time markers plus wish vocabulary, without an observable measure. No prioritisation, no must/teachable split.

What the workshop adds. The critical-few question ("Which three tasks decide in the first six months whether the hire holds?") leads to a reduction to eight levers instead of thirty. The hundred-percent distribution reveals a shift: the CEO gives the three most important activities sixty percent together, while the HR director rates these three at only forty percent, because in his view they overlook two further levers (onboarding architecture, reporting discipline). The triad leaves the difference in place and names it as a calibration point. For the success criteria, the moderation insists on observable measures per horizon: after three months "three growth hypotheses in writing, with owner and figure", after six months "one validated growth mechanism visible in the funnel", after twelve months "team of four people, own quarterly planning". Wish vocabulary becomes a picture of success.

Delta observation. The task list shrinks, the methodological depth grows. On his own, the CEO would have produced a broader list, but no measurable anchors. The workshop reduces the breadth and creates the depth.

6.2 Vignette B: HR director of a mid-sized company

Profile. HR director of a mid-sized company with eight hundred employees in the marketing/media/sales cluster. In 2024 the company faced a sharpening of its strategy and is hiring a new chief marketing officer. The HR director has handled five comparable searches in recent years.

Before the workshop. The HR director arrives with a list of fifteen requirements. The list is neatly ordered: four years of experience in a comparable position, business-fluent English, experience in B2C and B2B, leading teams of ten or more, a university of applied sciences degree, an understanding of balance sheets, brand experience. Every point carries equal weight. Asked which three of them are critical and which can be built up, she answers "all important". The persona of the ideal candidate has not been worked out; "I recognise the right person in the conversation".

What she would have produced alone. A complete list of requirements, neatly prioritised by years of experience and qualification stamps. A brief persona derived from the requirements ("someone who has done this before"). Clearly formulated search criteria for the market. But no must/teachable split, no persona why, no systematic check of which of the fifteen requirements can actually be operationalised in the selection process.

What the workshop adds. Separating must from teachable changes the picture. Of fifteen requirements, three remain in the must column: deep experience in aligning marketing and sales in a multi-tier sales model, leadership experience during a phase of organisational change, an understanding of the industry's brand architecture. Everything else - languages, degrees, industry knowledge in the narrow sense - moves into the teachable column or is dropped. The persona, previously implicit, is spelled out: "mid-forties, has been through a difficult brand sharpening in a similarly complex setup, is now looking for the role in which strategic sharpness and operational leadership come into balance". The candidate market shifts as a result: candidates who do not formally meet all fifteen points but hit the three must points and the persona become plausible.

Delta observation. On her own, the HR director would have produced the longer and tidier list, but also the more expensive search. The workshop shortens the list and changes the search profile.

6.3 Vignette C: head of operations at a manufacturer

Profile. Head of operations at a manufacturer with fifteen hundred employees, industry cluster operations. The position of head of a production unit is to be filled; the predecessor moved to another function after eight years. The head of operations knows the role at first hand, having held the position ten years ago.

Before the workshop. The head of operations arrives with operational detail. A task list of twenty-five tasks, ordered by weekly, monthly and quarterly rhythm. Three pain stickers on the current situation: a machine park that needs modernising, a shift leadership whose communication grinds, a supplier change in the background. Asked about the ninety-day plan, the answer is: "I have it in my head, it will fall into place". The persona of the person sought is very clear - "someone who wins the trust of the shift leaders in the first four weeks without positioning themselves as an outsider" - but documented nowhere.

What would have been produced alone. A rich current situation at pain-sticker level, which in our corpus is often the densest source. A complete task list that tips into overload, because every weekly rhythm adds further tasks. An implicit persona that does not appear in the brief and is therefore not available for the later selection. Success criteria as a time frame with wish vocabulary ("settled in after three months, productive after six months").

What the workshop adds. The hundred-percent distribution of tasks leads to the first point of consensus in the triad. The head of operations gives supporting the shift leadership thirty percent. The CEO gives it fifteen percent, because he takes it for granted. The HR director gives it twenty-five percent. The difference is

put on the table; the triad decides that this support explicitly counts as a core activity, because two of the three see it as critical. The persona is spelled out and documented: "has gained acceptance as external reinforcement in a similarly structured production unit without appearing as a stranger". The success criteria are pulled towards observable measures: after three months "shift handovers without escalation in two consecutive weeks", after six months "an own investment decision for the machine park submitted", after twelve months "the two following shifts staffed and briefed".

Delta observation. On their own, the head of operations would have produced the densest and most detailed self-reflection. But the persona would have stayed in their head, the triad difference would not have been disclosed, and the success criteria would have been phrased as wish vocabulary. The workshop discloses the triad consensus, documents the persona and moves the success criteria into observable measures.

> **The workshop is not the sum of the three views, but their friction.**

Convergence mechanics. The workshop is not the sum of the three views, but their friction. Where the CEO's strategy meets the HR requirements and the department's operations, the methodologically interesting layer arises. Where the three views automatically coincide (the description fields), the workshop is not the step that makes the difference. Where they diverge (the measurement fields), it is.

CHAPTER 07

Before-and-after metrics

We quantify the workshop delta in three dimensions. For each dimension, the pre-workshop anticipation (derived from building block a) is set against the post-workshop actual (measured from the forty-four boards). The quantification is conservative; where the data allows ranges, we show them.

D1: share of measurable success criteria per horizon

RATIONALE

Success criteria are the field in which wish vocabulary settles most easily. The question "What would be a success after six months?" can be answered without anyone insisting on observability. In two hours of self-reflection, stakeholders typically produce time markers plus wish phrases - that is the standard answer one's own head delivers. Methodological depth (one observable measure per horizon) only arises when the moderation enforces it.

FIGURE

Before the workshop, we estimate the share of measurable success criteria per horizon at close to zero. In the thirty boards captured from templates A and B (which mostly had open moderation), the post-workshop actual is also mixed: a median of eight notes per board (see synthesis report, finding 3 on success factors) masks the fact that most of these notes are time markers or wish vocabulary. Evidence in the corpus: B03 fills the field only with the three-six-twelve-month frame without measurable content; B25 fills it with wish phrases ("good collaboration", "meet internal deadlines"); B29 leaves the twelve-month marker empty. Positive evidence B04 shows that measurable pictures of success are possible when the moderation enforces the step ("get the current version running", "design a completely new app"). Robust estimate: in boards with moderation discipline (about twenty percent of the corpus), the share of measurable success criteria per horizon is sixty to eighty percent. In boards without this discipline (about eighty percent), it is below twenty percent. The delta between a high-quality workshop and self-reflection is thus roughly forty to sixty percentage points per horizon.

D2: share of prioritised key activities with a hundred-percent distribution

RATIONALE

Key activities is the field in which collecting becomes most disruptive. Without a stop rule, the field produces the longest list in the brief. A hundred-percent distribution of time or leverage across the activities is an analytical constraint that does not arise without moderation - it feels artificial and is left out in self-reflection. The workshop anchors it as a consistency check and uses it as a tool to make diverging stakeholder views visible.

FIGURE

Before the workshop, we likewise estimate the share of weighted key activities at close to zero. Self-reflection produces the collection, not the weighting. In the forty-four boards of the corpus, by our count, there are two clear cases of percentage weighting (B08 and B28), with B28 at the same time being the negative peak case with forty notes and no stop rule - the weighting was added afterwards. The median for the field is

eighteen notes; the peaks reach forty (B28), twenty-nine (B37), twenty-eight (B08), twenty-four (B36) and twenty-one (B07). For a post-workshop actual with a moderated stop rule and hundred-percent distribution, the corpus shows less than five percent. With a modern template (C and D), the distribution is anchored systematically; we estimate the share there at sixty to eighty percent. The delta between a disciplined workshop and an open collection is thus roughly fifty to seventy percentage points.

D3: share of requirements with a must/teachable split

RATIONALE

The must/teachable split is the most unambiguous methodological move the workshop produces. It forces the stakeholders to decide for each requirement: must this ability be present on day one, or can it be built up in the first year? This split is not intuitive. Self-reflection lumps; it produces the list without making the separation. The workshop splits, and exactly this split changes the search profile.

FIGURE

Before the workshop, we estimate the share of requirements with a must/teachable split at close to zero. In the thirty boards of template A, the list of requirements clumps together in a single field; the split does not exist. Only in templates C (four boards) and D (two boards) is the split anchored systematically. In the corpus captured, the share of boards with an explicit must/teachable split is therefore about fourteen percent (six of forty-four boards). With the more recent templates, this share would approach a hundred percent if applied consistently. The delta between a workshop in template A and a workshop in templates C or D is thus, measured by whether the separation actually exists, about ninety percentage points. That is the strongest template effect in our corpus.

CHAPTER 08

Recommendations

The findings lead to four lines of recommendation for different audiences.

8.1 Method - what Pre-Recruiting workshops need to do differently

Five interventions follow directly from the data. First: the must/teachable split as standard, whatever the template. The split is the most effective methodological move in the corpus and should not be reserved for the more recent templates. Second: a stop rule for key activities. Stop at around ten notes, group, enforce a hundred-percent distribution. Third: measurable success criteria per horizon. One observable measure for each three-six-twelve-month marker, not a wish phrase. Fourth: a fixed set of questions for Core Questions instead of free text. In our corpus, the open question systematically produces emptiness; the fixed set produces usable answers. Fifth: make a fundamental decision on Recruiting Process and Core Questions as fields. Both fall below the fill-level threshold at which a field can be considered load-bearing. Either they are redefined or they are dropped.

External triangulation: existing recruiting benchmarking reports have systematically captured outcome metrics such as time to fill, cost per hire and quality of hire for years (SHRM Recruiting Benchmarking 2025, https://www.shrm.org/topics-tools/research/2025-recruiting-benchmarking). What they do not measure is the quality of the briefing layer that determines these outcome metrics upstream. Our study closes this gap at the source, not at the result.

8.2 Moderation - what moderators need to ask differently

Three interventions in moderation. First: always lead the success-factors question to an observable measure per horizon. The question is not "What would be a success?", but "How will the team recognise success after six months, in a measure that two observers could establish independently of each other?". In the data, this question reproducibly produces better answers. Second: use the hundred-percent distribution of key activities as a consistency check for the triad. Not as a technical step, but as a tool to make the diverging views of the three roles visible. Third: metadata as a mandatory checklist step. In our corpus, boards without a date, briefing partner or job title are not the exception - that is a lack of discipline that moderation can resolve.

8.3 Content - libraries, persona, pain typology

Three lines of content. First: reusable requirement libraries per job family. The patterns repeat strongly in our corpus - tech roles with SAFe and agile, marketing roles with performance and funnel, leadership roles with change and scaling. A library per family reduces the collecting work and makes room for the methodological layer. Second: anchor persona and why early. In template A, both are under-represented; in template D, the gap closes. In the corpus, the persona is often the department head's strongest point and management's weakest; the workshop should force the switch between perspectives early. Third: use the pain typology (predecessor, interim, reorganisation, growth) as a conversation framework for the current situation. It is not exhaustive, but it asks the stakeholders about the starting point without forcing answers.

8.4 Implication for the Pre-Recruiting industry

Pre-Recruiting is not "starting earlier". Pre-Recruiting is "starting in a methodically measurable way". The thesis is not meant polemically. It is what the data shows: the workshop effect does not lie in the depth of

collection, but in the methodological layer. Anyone selling Pre-Recruiting is either selling that layer or selling an attractive methodological coat over a collection. The difference can be read from the workshop artefacts. One consequence for the industry: the discussion about Pre-Recruiting should not stop at the level of "we do that too", but move to the level of "we make it measurable at the following three points". These are exactly the three points this study has identified.

The Leadership IQ study "Why New Hires Fail" (Mark Murphy, 5,247 hiring managers in 312 organisations, more than 20,000 hires recorded) is the quantitatively most robust outcome study on this point: forty-six percent of newly hired employees fail within eighteen months. Only eleven percent of them because of a lack of technical skills. The remaining eighty-nine percent fail because of coachability, emotional intelligence, motivation and temperament (https://www.leadershipiq.com/blogs/leadershipiq/35354241-why-new-hires-fail-emotional-intelligence-vs-skills). This outcome distribution matches our briefing observation: exactly the factors on which hires later fail are the factors that remain structurally unaddressed in template A. The must/teachable split, the persona-why question and measurable success criteria address precisely this layer. The template evolution in our corpus is methodologically consistent with outcome research outside it.

CHAPTER 09

Appendix

A1. Template distribution

TEMPLATE	NAME	BOARDS
A	HireWorks recruiting canvas (1 to 9)	30
B	ThinkHeads recruiting canvas	8
C	Hiring Requirements Engineering	4
D	High-Impact Hiring Blueprint	2

Atypical boards without the standard grid were classified separately and are not included in the bucket evaluation. Seventeen of forty-four boards (about thirty-nine percent) have a separate "Welcome" screen in the export.

A2. Anonymisation protocol

Each board was given a B-number (B01 to B44) and an industry cluster in the master. Real names, personal names and specific locations remain in the internal master. Public statements in this study use only the B-ID plus industry cluster. The stakeholder vignettes in chapter 6 are generic archetypes; they are placed in an industry cluster but are not reconstructed from a specific individual board.

A3. Methodology note

Extraction: forty-four boards verbatim (B01 to B04 directly, B05 to B44 via eight parallel subagents with a strict verbatim and anti-fabrication protocol). Quality check: B25 (template C, industry cluster finance) checked field by field against the original image, matching exactly including the colour coding of the sticky notes. Counting: note counts per common bucket determined programmatically, with a parsing gap for CF1 to CF4 for about thirteen boards whose note lines were formatted differently (raw data complete in the master). Saturation: inference, not a formal code count per note.

A4. Limitations

Three limitations should be kept in mind when reading. First: the corpus leans towards digital and tech. Classic industrial key roles are under-represented; the findings are calibrated to knowledge-intensive leadership and specialist roles. Second: template heterogeneity makes longitudinal statements difficult. We compare fields across template generations only where the bucket logic holds. Third: the saturation point is inference, not a formal evaluation. A larger study with a formal code count per note would be technically possible but methodologically no more productive; the dominant categories will not shift.

A5. Glossary

Pre-Recruiting. The methodologically invisible phase that precedes a key hire. Briefing, triad consensus, sharpening of requirements, calibration of success. Pre-Recruiting is the category in which HIHB works as a method.

HIHB (High-Impact Hiring Blueprint). The method documented in this study. It integrates persona, why, impact scoring and a must/nice split and is the most recent generation of the HireWorks workshop templates.

5C Method. The toolset HIHB uses in moderation. The five steps are Contingency, Consistency, Calibration, Coordination, Clarification. In this study, 5C appears indirectly in three places: the consistency logic in the hundred-percent distribution of key activities, the calibration logic in the must/teachable split, the coordination logic in the triad.

Critical few. The selection rule that filters the few load-bearing activities out of a broad list. Put into practice through a stop rule and a hundred-percent distribution.

Triad. The three stakeholder roles that typically come together in a key-hire workshop: management, HR, department head. The triad consensus marker documents where the three views meet and where they diverge.

Brief. The workshop artefact that methodically carries the key hire. The subject of this study.

Workshop delta. The difference between what stakeholders produced AFTER two hours of moderated workshop and what they would have produced WITHOUT a workshop in two hours of self-reflection. The term sums up what this study measures.

Must/teachable split. The separation in the list of requirements between abilities that must be present on day one and abilities that can be built up in the first year. First anchored systematically in template C.

This study is part of the HIHB research series. The underlying data (forty-four boards, fill-level matrix, template census) is held in the internal master and is available on request in anonymised form for methodological queries. Contact: HireWorks GmbH, Hamburg. This is a translation of the German original.

HIHB

WHAT TO DO NOW

15 minutes are enough to ask the right question.

If you are currently filling a key role, the next sensible move is not another briefing document, but a moderated workshop that operationalises the three measurement fields of this study. In 15 minutes we look at your role together and decide whether a HIHB workshop is the right next step. Free of charge, no obligation.

Book a 15-min fit call -> calendar.app.google/5tnSKPHybKAThwzg9