

INFRASTRUCTURE RESEARCH

THE OPERATIONAL VISIBILITY GAP

Why recruiting breaks at scale.

A research report documenting the structural failure modes of fragmented hiring execution — and what genuine operational infrastructure requires. For senior talent leaders, recruiting operations, and hiring executives.

68%

VP searches stall past week 10 due to execution failure

2024–25 VP-level mandate observations, internal Majhi OS recruiting operations

14 wk

Avg time-to-fill on fragmented tooling

Median across tracked VP+ mandates with fragmented stack

\$3.8K

Monthly tool spend per recruiter before first hire

Composite estimate across 6 typical vendor categories

50 d

Avg close time on unified infrastructure

Majhi OS 2024 engagements, VP-level mandates



EXECUTIVE SUMMARY · SECTION 01

The problem isn't talent scarcity.

01

A VP Engineering search enters week 9.

The recruiter believes the pipeline is healthy. The hiring manager believes outreach is still active. Leadership assumes finalists are already in process.

IN REALITY:

- Reply rates collapsed 3 weeks ago — undetected, because no alert threshold was set
- Two candidates were duplicated across mandates — both received identical outreach sequences
- No recruiter activity occurred in the past 96 hours — invisible in the current tool stack
- The outreach domain reputation degraded silently — future sequences are now pre-penalized

No one knows. Because the recruiting stack cannot see itself.

This is not an edge case. It is the median experience of a VP-level search on fragmented tooling. **The primary cause is operational invisibility** — not market conditions, not candidate scarcity, not recruiter capability. The tools simply do not share state, produce no unified signal, and offer no systemic view of what is happening inside a search.




Stall Rate

VP searches stalling past week 10 on fragmented tooling



TtF Reduction

Faster time-to-fill vs. fragmented average (VP level)



Audit Coverage

Fragmented teams with full timestamped decision trails



The difference between a search that closes in 50 days and one that closes in 120 is rarely the talent market. It's almost always the execution layer — and whether anyone can see inside it.

— Manas Majhi, Founder, Majhi OS

SECTION 02 · BENCHMARK DATA

The real cost of fragmented hiring

02

STRATEGIC CONTEXT

The cost of fragmented hiring is not software spend. It is **execution delay** — the gap between when a mandate should close and when it actually does. At the VP level, every week of delay carries compounding cost: leadership bandwidth, candidate attrition, opportunity cost. That is the metric that matters.

Time-to-Fill: VP-Level Role Comparison

Median weeks across mandates. Based on 2024–25 internal Majhi OS mandate observations.




Each week of delay compounds: candidate attrition accelerates, market windows close, leadership attention fragments. Operational compression — not tooling speed — is the lever.

Execution Fragmentation: Tooling Overhead


Monthly per-recruiter tool spend on a typical fragmented stack vs. a unified operational system.

Fragmented Stack
\$3,280/mo



Unified Infrastructure Layer

1 operational system · 1 source of truth
Reduced workflow fragmentation
Reduced coordination overhead



SECTION 03 · ROOT CAUSE ANALYSIS

Four structural failure modes — and their consequences

03

01

Intake without normalization

Candidates arrive from LinkedIn exports, CSV uploads, referrals, and ATS imports in different formats with no deduplication. The same candidate may enter the funnel twice with no system to detect the collision.

OPERATIONAL CONSEQUENCE

- › Quality gates fail silently on unstandardized profiles
- › Duplicates consume recruiter capacity without visible cost
- › Mandate timelines extend without a traceable cause

02

Outreach without verification

Email sequences run on unverified contact data without DNS/MX verification or bounce-kill logic. Domain reputation degrades silently with each undeliverable message.

OPERATIONAL CONSEQUENCE

- › Domain reputation degrades — future sequences pre-penalized
- › Reply rates collapse progressively without visible root cause
- › Engagement failure discovered retrospectively, weeks late

03

Evaluation without evidence

Hiring recommendations are made on résumés and recruiter judgment — none of which is systematically recorded or auditable. When a board asks “why this candidate?”, the answer cannot be defended.

OPERATIONAL CONSEQUENCE

- › Shortlist rejection by board or CEO — rework cycles initiated
- › Recruiter credibility erodes with hiring managers over time
- › No audit trail if a challenged hire is escalated

04


Execution without observability

A search can be technically in progress — messages sending, profiles being reviewed — while the pipeline is functionally stalled. Without real-time mandate health states and stall detection, failure is invisible until it's already expensive.

OPERATIONAL CONSEQUENCE

- › Searches stall undetected for 2–4 weeks before escalation
- › Leadership receives delayed status signals — trust degrades
- › Mandate recovery begins too late to avoid timeline impact

The Compounding Failure Loop — how each mode amplifies the next



Each cycle of manual recovery increases technical debt in the search — and the next mandate starts from the same broken foundation.

SECTION 04 · MATURITY MODEL

Fragmented operation vs. infrastructure operation

04

The gap between these two operating models is architectural — not a matter of tooling budget or headcount. The question is whether execution generates a unified signal across stages, or whether each stage operates as an independent silo.

DIMENSION	✗ FRAGMENTED OPERATION	✓ INFRASTRUCTURE (MAJHI OS)
Intake	Multiple formats, no normalization, manual deduplication	Normalized intake with duplicate detection and profile stitching
Evaluation	Résumé + gut feel, no structured evidence or mandate scoring	Evidence dossiers with proof points, risk flags, and fit scores
Outreach	Unverified data, manual follow-up, no bounce logic	DNS/MX verified, bounce-kill active, stop-on-reply automated
Pipeline	Spreadsheet + weekly calls; problems surface 2–3 weeks late	Live mandate health states with funnel decay alerts and stall detection
Reporting	Manual exports, always retrospective, no real-time signal	Real-time telemetry across all active mandates in one command surface
Auditability	No record of actions; decisions untraceable if challenged	Full audit log — every event, action, and override timestamped
Human Governance	Manual coordination; no structured override or escalation record	Structured override workflows with full event logging; recruiter approvals at every gate; escalation paths defined and auditable

Operational note: Mandate health visibility reduces delay detection time, but recruiter intervention remains required during market volatility, low-signal searches, or when candidate availability degrades mid-mandate. Infrastructure improves signal quality — it does not remove the need for experienced judgement.

PERFORMANCE BENCHMARKS


Key Operational Metrics: Fragmented vs. Majhi OS

Based on VP-level mandates. 2024–25 observations.

Fragmented Stack Majhi OS

Time-to-Fill (VP level)

Median weeks



Outreach Reply Rate

With vs. without DNS/MX verification




Shortlist Approval Rate

Board / CEO accepted shortlists (evidence-backed vs. not)



Full Audit Trail Coverage

% of mandates with timestamped end-to-end decision log



SECTION 05 · PIPELINE ANALYSIS

Where candidates are lost in the funnel

05

The funnel comparison illustrates how each failure mode compounds at consecutive stages. The percentages shown are directional — actual results vary significantly by role complexity, market conditions, and recruiter maturity. The structural pattern, however, is consistent: fragmented tooling creates attrition that gated execution prevents.

Recruitment Funnel: Fragmented Stack vs. Majhi OS

Illustrative per 400 sourced candidates. Median VP-level mandate. Results vary by market and role complexity.

STAGE	FRAGMENTED STACK	MAJHI OS	Δ
Sourced	400	400	—
Normalized / QMS	210	380	+81%
Dossier Ready	160	365	+128%
Outreach Sent	140	340	+143%
Response Received	48	160	+233%
Interview Stage	24	90	+275%
Hire	8	52	+550%

Methodology: Illustrative median funnel across selected VP-level mandates. Fragmented column reflects unmanaged attrition on a typical 6-tool stack. Majhi OS column reflects gated, evidence-backed execution with QMS normalization and verified outreach. Results vary significantly by market, recruiter maturity, role complexity, and search timeline. These figures represent directional patterns, not guaranteed outcomes.

WHY THE DIVERGENCE COMPOUNDS

The dramatic funnel divergence is not caused by a single feature. It accumulates across every gated handoff: normalized intake retains 95% vs. 52% into QMS; verified outreach delivers 3× reply rate; evidence-backed shortlists produce 2× approval rate. The infrastructure advantage is architectural — it compounds at every stage, not just one.



SECTION 05B · EXECUTIVE IMPACT

Operational signals. Business consequences.

05

Recruiting delays are rarely isolated workflow problems. They compound into leadership drag, execution risk, and organizational slowdown. Every infrastructure signal Majhi OS surfaces maps to a direct executive business outcome.

OPERATIONAL SIGNAL	METRIC	EXECUTIVE BUSINESS OUTCOME
Mandate compression	14 wk → 7 wk	Faster executive hiring velocity. Leadership roles filled at half the industry median — compressing time-to-contribution and reducing opportunity cost.
Reply-rate recovery	14% → 35%	Improved pipeline continuity. DNS/MX-verified outreach reaches candidates fragmented stacks miss — keeping pipelines alive without manual intervention.
Real-time mandate visibility	Live telemetry	Reduced leadership escalation. Hiring managers receive status automatically — no weekly sync, no surprises at week 10.
Stall detection	2–4 wk earlier	Earlier intervention before failure. Automated decay alerts surface degradation before it becomes a missed deadline.
Full audit trails	100% coverage	Board-level defensibility. Every candidate decision timestamped and evidence-backed. When the board asks why — the answer is ready.
Shortlist approval rate	38% → 82%	Fewer rework cycles. Evidence-backed dossiers reach boards ready to approve — cutting shortlist-rejection loops that stall searches by weeks.

Majhi OS improves not just recruiting workflows — but the operational reliability of leadership hiring itself.

Recruiting delays are rarely isolated problems. They compound into leadership drag, execution risk, and organizational slowdown. Controlled execution — backed by infrastructure — is the only kind that scales without breaking.

SECTION 06 · ARCHITECTURE · PART I

System boundaries and the execution loop

06




PRODUCT BOUNDARY


Majhi OS is the operational orchestration layer for high-stakes recruiting workflows — sitting between fragmented sourcing inputs and downstream hiring systems, creating unified visibility, evidence, and control.

System Boundary: Where Majhi OS Operates

INPUT LAYER — FRAGMENTED SOURCES



MAJHI OS — OPERATIONAL ORCHESTRATION LAYER



OUTPUT LAYER — HIRING STAKEHOLDERS



CORE EXECUTION LOOP

```

majhi-os · execution-loop · mandate-0041 · vp-engineering

01 Candidate PDF uploaded · intake normalized · format standardized · freshness verified
02 Identity stitched across LinkedIn + CRM + mandates 0038/0039 · no duplicate collision
- [ ALERT ] Duplicate candidate cluster detected: candidate-0291 appears in mandates 0039 and 0041
- [ HUMAN ] Manual deduplication · mandate 0039 record retained · 0041 entry archived · reason logged

03 QMS gate · fit score 8.7/10 · mandate alignment: STRONG · evidence confidence: HIGH
04 Dossier generated · 7 proof points · 2 risk flags surfaced · verification state: confirmed
05 [ HUMAN ] Recruiter review · risk flags acknowledged · dossier approved · outreach unlocked

06 3-touch sequence queued · DNS/MX verified · bounce-kill active · stop-on-reply enabled
- [ ALERT ] Reply rate -28% vs. mandate baseline · stall probability: 74% · escalation queued
07 [ HUMAN ] Manual override · sequence paused · reason logged · alternative approach initiated
- [ ALERT ] Mandate health: AT RISK · leadership dashboard updated · telemetry event written
- [ HUMAN ] Mandate reprioritization initiated · recruiter reassignment logged · stakeholder notified
  
```

SECTION 06 · ARCHITECTURE · PART II

Failure recovery logic and human governance

06

Infrastructure earns trust not by appearing frictionless, but by behaving predictably under stress. Real recruiting operations encounter ambiguity, degraded states, and conflicting signals. The system's value is in how it responds when things go wrong — not just when they go right.

```
majhi-os · failure-recovery · mandate-0041 · degraded-state
```

- [ALERT] Funnel decay threshold exceeded · reply rate: 6% (baseline: 28%) · 18-day window
- [ALERT] Recruiter overload flag · 4 active mandates above capacity threshold
- Recovery sequence initiated automatically · outreach paused across all active touches
- [HUMAN] Escalation acknowledged · root cause analysis initiated · notes logged to audit record
- [HUMAN] Mandate health downgraded: ACTIVE → AT RISK · leadership visibility enabled
- [HUMAN] Alternative sourcing channel queued · sequence logic revised · reason captured
- Recovery validation · reply rate recovering: 6% → 19% (day 8) · mandate restored to ACTIVE
- Full audit record written · trigger + action + outcome + timestamp + recruiter ID · immutable

Recovery Sequence: What Infrastructure Does When a Mandate Degrades

- SYS Automatic detection:** Outreach paused when reply rate falls below mandate baseline threshold. No manual monitoring required.
- SYS Escalation routing:** Recruiter escalation initiated with pre-populated context — timeline, trigger event, current pipeline state.
- HUM Human governance:** Mandate health downgraded manually after recruiter review. Reason logged. Override captured in audit trail.
- SYS Leadership visibility:** Mandate health dashboard updated automatically. No status update call required. Signal propagates in real time.
- HUM Recovery approval:** Alternative sequence approved by recruiter before reactivation. Confidence range attached. Rationale documented.
- LOG Audit record:** Full event log written — trigger, action, override, outcome, timestamp, recruiter ID. Immutable. Board-ready.

THE AUDITABILITY TEST

For any hiring decision: can you show the exact sequence of events that led to this recommendation? Can you show which risk flags were surfaced and what was done about them? Can you show who approved the override and when? If the answer to any of these is no — you are one board question away from a credibility problem.

SECTION 07 · DESIGN PRINCIPLES

Six principles for infrastructure-grade hiring

07

These are architectural commitments — decisions made at the design level that determine whether a recruiting operation can scale without breaking. Each principle is grounded in observable operational behaviour, not philosophy.

PRINCIPLE 01

Evidence before opinion.

Every recommendation must be traceable to structured evidence — proof points, fit scores, risk flags. Unrecorded judgment is a liability at board level.

In practice: When a hiring manager asks why a candidate was shortlisted, the recruiter opens a dossier — not a memory.

PRINCIPLE 02

Observability is foundational.

If you can't see what's happening inside a search in real time, you can't manage it. The command center is the primary control surface — not a reporting tool.

In practice: The system detects declining reply rates before leadership asks why the mandate stalled.

PRINCIPLE 03

Gate before you execute.

Every candidate passes through a quality gate before outreach. The QMS is what makes scale trustworthy, not just fast.

In practice: A candidate with a fit score below threshold cannot enter an active sequence without explicit recruiter override — logged and auditable.

PRINCIPLE 04

One system. No handoffs.

Every tool handoff loses context and slows execution. Infrastructure means one connected system where data flows unbroken from intake to close.

In practice: A candidate record updated in intake is immediately visible in outreach, evaluation, and leadership reporting — with no sync delay.

PRINCIPLE 05

Automation reduces load.

The test: does it give the recruiter more clarity, or more to manage? Complexity without clarity is a more expensive version of the problem.

In practice: Bounce-kill and stop-on-reply run automatically. The recruiter never manages sequence state — only approves escalations.

PRINCIPLE 06

Surface failures first.

Real infrastructure surfaces problems before humans discover them. Stall detection and decay alerts are the primary signal the system is working.

In practice: A mandate health downgrade triggers an automatic escalation — the recruiter receives context before the hiring manager calls.

A NOTE ON HUMAN GOVERNANCE

None of these principles reduce recruiter accountability — they increase it. Structured override workflows, approval gates, and confidence ranges don't automate decisions. They make decisions traceable. The recruiter remains responsible for every outcome. The system makes that responsibility visible.

CONCLUSION

Infrastructure thinking, not tool thinking.

08

The recruiting industry has invested heavily in tools that make individual tasks faster — better search, better enrichment, better sequencing. These investments are not wasted. But they do not close the operational visibility gap, because the gap is not caused by slow individual tasks.

The gap is caused by the absence of a connected system that produces a unified signal across all tasks simultaneously. Closing it requires a different architectural decision: to treat recruiting execution as an operational system, not a collection of workflows. To build around observability. To make every candidate event, recruiter action, and quality gate decision visible, logged, and auditable.

Controlled execution is the only kind that scales without breaking. That is the commitment Majhi OS was built to keep.

50_d

Average VP close time on Majhi OS (2024)

↓ 50% vs. fragmented avg

35%

Outreach reply rate with DNS/MX verification

↑ 2.5x unverified baseline

82%

Shortlist approval rate — board and CEO

↑ from 38% unstructured

100%

Mandates with full timestamped audit trail

Every decision logged



REQUEST AN OPERATIONAL WALKTHROUGH

Explore the command surface, event telemetry, and evidence-backed execution layer using your own hiring workflows. A Mission Walkthrough is 45 minutes with a senior Majhi OS team member — using a live mandate brief, your context, your role complexity.

majhi.tech · support@majhi.tech

Methodology & attribution: All benchmark data is derived from Majhi OS internal mandate observations (2024–25) and publicly available research on executive recruiting timelines. VP-level is defined as VP, SVP, EVP, and C-suite. Funnel data is illustrative, based on median search patterns. Results vary by role, market, and team configuration. Pricing comparisons use composite estimates from publicly available vendor pricing at standard seat counts — no direct vendor data was used. This document is intended for executive and research audiences and does not constitute a commercial offer.