RPM - AI Search Report

Project: AI Search on RPM
Test Period: May 2026
Author: Yuhan Hu (UXD)
Sessions: 7 sessions, 8 participants, across Recruiter, RC, and Sourcer roles
Status: Round 1 complete. Additional sessions scheduled with high-volume recruiters and worklist users.

Executive Summary

We tested an interactive prototype of the AI search feature on the Hire Recruiting Pipeline page with 8 participants across recruiter, recruiting coordinator, and sourcer roles. The test focused on three questions:

1. How do users react to the AI search experience?
2. What content do different roles need to search for?
3. What additional functional needs surface around the workflow that AI search touches?

Headline findings:

- Reception is unambiguously positive. Every participant validated the core flow with strong language ("really cool", "100% would utilize", "very intuitive", "miles better"). One recruiter, Jon, said he would stop using filters entirely once AI search ships.
- The biggest design success is the pattern itself. Natural language input, AI-generated filter chips distinguished from system filters, partial result handling, and the deliberate "search bar, not chatbot" framing for P0 all landed without friction.
- Different roles want fundamentally different content from AI search.
    - RCs want SLA-aware queries and event status.
    - Recruiters want cross-view task management and HM context. Sourcers want re-engagement and notes-based queries.
- The most consistent unmet need is **data reach**. Across 5 of 7 sessions, participants asked AI search to reach into data that lives outside RPM today: candidate profile notes, C3 messaging, scheduling history, etc. AI search exposes this gap; it does not create it.
- AI search suggestions (“Try search for”) are not one-size-fits-all. Three different mental models surfaced (recent history, pinned commonality, tenure-based fade-out), and nobody asked for the same suggestion content.

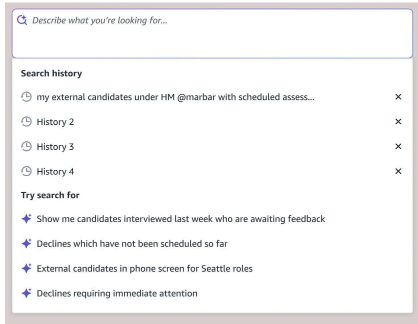

Recommendations

1. Natural-language input plus AI-generated filter chips with color distinction.
2. Result summary line plus chip row. The duplication is informational, not redundant. Camille validated this directly.
3. Partial-result handling with clear "we couldn't find X" messaging.
4. Disambiguation prompts for ambiguous IDs and names.
5. **A static suggestion list at launch (need TA leads to provide examples), four examples or fewer, with an explicit plan to personalize after one to two months of usage data.**
6. **Search history of 4 recent items, feasibility-pending.**
7. **Type-ahead from history. Highly valuable per the Madhurima's team and Emily. Ship if feasible.**
8. Search-bar-not-chatbot framing. Keep as-is.
9. The rollout plan should include training and framing language so the team understands AI search is iterative, not magic on first try.

Participants (Optional Background)

- Camille, Recruiting Coordinator (Tech). LA hub, Mexico City focus. SDE / SDM scheduling.
- Madhurima, RC Manager / Process Lead, NAMER region. Tested jointly with Ubair.
- Ubair, RC / Scheduler / SME, NAMER region. 4 years on GTS, L1 to L7 loops.
- Jon Ormond, Senior Full Life Cycle Recruiter. EMEA Legal. Team tech POC.
- Kathy Ziff, L6 Senior Specialty Recruiter, GSR L7+ scientists. Bar Raiser trainee, GSR tech advisor.
- Michael, Technical Sourcer (grouped with recruiters in this report). Audible Tech. Full-life-cycle behavior despite sourcer title.
- Ella, RC Lead, GSR Executive team. Works closely with the UX team.
- Emily, RC and QPOC, AIGC plus Entertainment Fit. QPOC for on-site distribution. Roughly 5 years tenure.

Coverage was strong on RC and Recruiter, with one Sourcer. The gap is Worklist users because they don't have experiences with current RPM yet.

Reactions to Current AI Search Design

1 Overall Sentiment

Sentiment was positive across all sessions.

Camille said "that's really cool, 100% would utilize." Madhurima and Ubair said it was "very intuitive" and "pretty much excited to put this into production." Jon said it would be "answering off the top of my head kind of all of the issues I would currently see." Kathy said "oh, I love that" when she saw the AI-generated filter chips. Michael said "nice, that is so helpful." Ella said "I like how simple it is and looks on the screen too." Emily said it was "gonna make a really big difference."

The strongest single signal came from Jon, who explicitly said he would stop using the filter UI entirely once AI search launches.

2 What Worked in the Design

- Natural language input had strong positive reception across all 8 participants. Spelling tolerance was specifically called out.
- The visual distinction between AI-generated filter chips (purple) and system filters (gray**)** was understood without explanation. Kathy's reaction was the most enthusiastic.
- The result summary line ("Found 18 candidates...") was received positively. Madhurima noted that it "tells me how much of my work is left for the day."

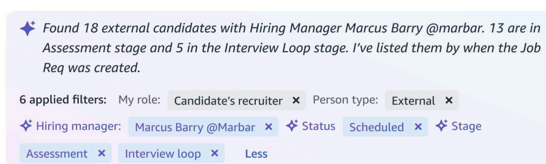

- The duplicate filter display, where filters appear both in the AI summary line and as removable chips, was acceptable to positive. Camille said it actually helped her understand what the AI extracted.
- Partial result handling, the "we found these but couldn't find ice cream" pattern, was endorsed by all participants. Jon called it "the logical way of doing it rather than doing nothing."

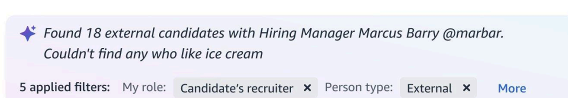

- The disambiguation pattern, where typing "Justin G" returns a list of Justins to pick from, was received positively. Multiple users said the current type-switching dropdown is painful.

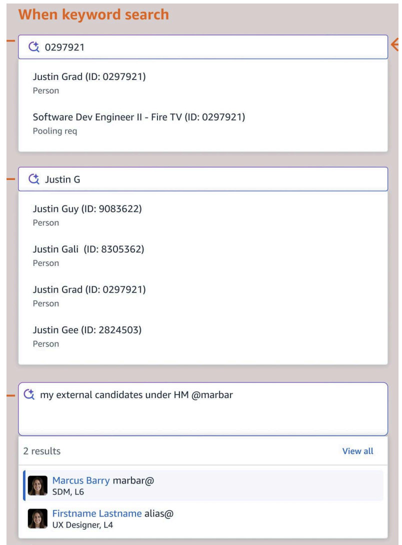

- The "search bar, not chatbot" framing for P0 was endorsed across the board. All participants understood why we are starting conservatively.

3 Friction Points and Pushback

- **Suggestion content felt generic.** The designer noted in-session that no two participants asked for the same suggestion content. Camille wanted more variety, not quantity, to discover capability boundaries. Jon would deprecate suggestions entirely for tenured users (3+ months). The recommendation is to plan for personalized, role-aware suggestions, ship a static list at launch, but design the system to evolve.

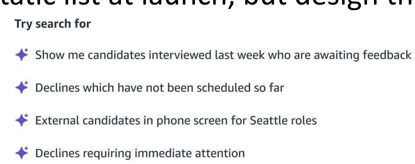

- **Search history count.** Three different opinions surfaced. Jon said 4 is optimal, with a hard ceiling of 8. Madhurima preferred 6. Camille had no specific preference but supports the concept. Kathy preferred **pinned commonality** over recent history (Asa-style pattern). The recommendation is to start with 4 recent items and revisit pinned or personalized later.

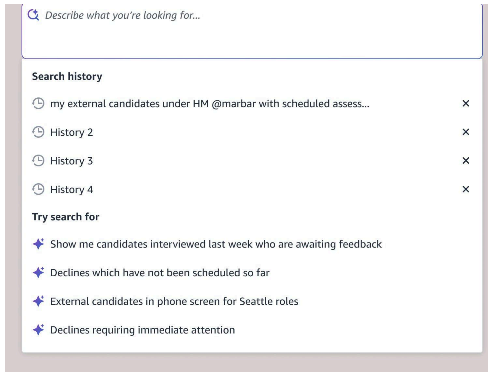

- **Where the search lives.** Kathy wanted AI search to span all **views**, not be scoped to a single tab. Full cross-view aggregation is likely an AI agent capability for Q3 and beyond.

4 What Users Did Not Misunderstand

The purple-color signal for AI-generated filters, the interpretation summary as the AI's translation of the query, the chips as removable filters, and the conscious decision not to be a chatbot in P0 were all understood without prompting. These are confidence-building signals: the design pattern itself is solid.

Role-Based Content Needs from AI Search

This section is the content backlog. Different roles need different things, and designing one universal suggestion list would underserve all of them.

1 Recruiting Coordinators (Camille, Madhurima, Ubair, Ella, Emily)

- SLA and time-sensitive queries. Candidates pending 2nd contact, pending 3rd contact, requests close to past SLA, 1st contact idle for 18 hours or more.
- Event status queries. Last-24-hour declines (Madhurima), no-show triage (Camille), upcoming debriefs missing a Bar Raiser (Emily), repetitive interviewer declines (Emily).
- Logistics queries. In-person versus virtual interviews (Ella, Emily), booked-room status (Camille), debriefs scheduled in the next 48 hours with unconfirmed interviewers.
- Cross-recruiter triage. "Pull up Ted's candidates" (Camille), "all loops under one recruiter" (Emily). RCs are aligned to specific HMs, so HM-based search has low filtering value for them.
- **Notes-driven** queries from request notes. By Bar Raiser, by interviewer, by location. These all live inside request notes today (Ella). Note templates
- Per-candidate aggregation. Show me everything about a candidate, including pre-brief, interview, and debrief, scheduled or not (Ella).
- Interviewer calibration lookup. Interviewer's calibration status or on-site count (Emily, currently done via the AI keyword search bar).
- Delegation scenarios. When taking over a teammate's queue during sick leave, search across the delegated profile (Madhurima).
- Cross-system data. Candidate availability confirmation status, which lives in C3 and Outlook today, not RPM (Camille).

- A note specific to RCs: pending feedback is not a relevant status for them. Ella confirmed it was removed for RC roles. Suggestions for RCs should not include this.

2 Recruiters (Jon, Kathy)

- HM-context queries. "How many candidates do I have at phone screen and loop for [HM]?" Kathy uses this for HM prep meetings.
- Search Health Index reporting. Per-rec breakdown such as "Rec 1: 2 phone screens, 1 loop. Rec 2..." Kathy uses this for a weekly report.
- Cross-view task management. Candidates I'm interviewing in next 48 hours, plus interviews still pending feedback. This needs to span recruitment queue and candidates-I'm-interviewing (Kathy, with her dual recruiter and BR role).
- Forward-looking summaries. Summary of all interviews next week, flagging declines and missing interviewers (Jon).
- Backward-looking compliance. Candidates interviewed last week, awaiting feedback.
- HM-multi-rec scenarios. All applicants for roles under a given HM, when one HM owns multiple recs (Jon).
- Cross-system data. C3 messaging history queries like "who haven't I responded to?" and "whose messages haven't been opened?" (Kathy).
- GSR pooling-rec hygiene. Slice through generic L7 / L8 / L10 pooling recs to find personally owned activity (Kathy).
- A note specific to recruiters: they care more about event timing and cross-view task triage than about scheduling logistics. Default sort for them should favor activity date over rec creation date.

3 Sourcer + Recruiter (Michael)

- Sourcer is grouped with recruiters in this report (Michael does end-to-end work despite the title), but the search content is genuinely different.
- Pipeline triage. Candidates requested for interview but not 100% confirmed by all interviewers.
- Re-engagement and silver medalist queries. "Candidates interviewed 2 years ago with a recommendation for recycle." "Mixed but close-to-inclined feedback recommended in 2 years."
- Prospect filtering, resume-style. NJ-based, not from a bank, lead software developer with scalable architecture experience.
- Notes-based search across candidate profile notes. Candidates where I noted "holding on" or "great SDE2 when available."
- Prep tracking. Have I had the prep call with this candidate? Candidates waiting on me to follow up because they haven't provided availability.
- Debrief intelligence. Find an incline and a decline debrief for SDE 3 in consumer domains, with detailed reasons. For prep reuse.
- Org and rec context. A 3-bullet snapshot of what each rec actually expects, since different niches can sit under the same title.

- A note specific to sourcers: many of Michael's most valuable queries (re-engagement, debrief intelligence, prospect filtering) likely belong to the parallel Sourcing AI Agent project. The boundary needs to be confirmed before P1.

4 Cross-Role Convergence

Some search themes appeared across all three roles, even if framed differently.

- Action-required queries (the "what needs my attention" framing) were universal. This is the strongest framing for default suggestions across all roles.
- Forward- and backward-looking time queries were universal. Both "what's coming up" and "what did I miss" are core.
- Cross-system data reach was raised by every role. Each surfaced a system AI search needs to eventually read from.
- Event-level status data, including declines, no-shows (?), confirmations, and feedback overdue, came up across all roles. Most of it sits at the data layer below current RPM.

Open Questions for PM and Engineering

1. Notes searchability. Can AI search read structured or templated request notes for P0? Raised by Ella, pending engineering input. [Template]
2. C3 and cross-system reach. Is it possible for AI search to query across systems?
3. Last-24-hour decline detection. Is the timestamp accessible from RPM, or only from the scheduling page?
4. Referral source data. Can the data be re-exposed in Hire?
5. Personalized suggestions. What infrastructure is needed for AI to learn from per-user query history?

Outstanding Research Gaps

Round 2 of testing should target: AUTA users, who are moving onto RPM in the future.

Additional Feature Needs Beyond AI Search

These came up in open follow-up questions ("anything else you'd like to share about Hire?"). They are not AI search asks, but they are real pain points the AI search effort surfaced.

1 Data Visibility Gaps in the Pipeline UI

- Add HM info to the job rec filter view. The current view shows rec ID and title only, which is unusable for Jon (his team has 3 to 4 job titles spread across many recs). This is the top recruiter pain point in EMEA Legal. Quick fix, high impact.
- Restore referral source visibility, which was lost when migrating from ICIMS. Jon says this is the most frequently asked of any item he hears as team tech POC.
- Booked-room visibility on the pipeline main view, so RCs don't have to click into each phone screen (Camille).

2 Workflow and Cross-Page Asks

- No-show candidate triage flow. Camille called this the "most frequent RC emergency." Today it requires scrambling across Outlook, Asana, Hire, and the interviewer in roughly 30 seconds. Aligns directly with the Q3 events status project.
- Per-candidate aggregated view across request and scheduled (pre-brief, interview, debrief). Raised by Ella as a search need but also valid as a UI feature regardless of AI.
- Repetitive interviewer decline flagging, so RCs can escalate to the HM. Raised by Emily, mirrored by Madhurima and Ubair. RCs role-specific.
- Phone-screen-to-on-site flow on filled jobs. Block before it reaches the RC team queue. Raised by Emily. Already addressed by the in-flight job-detail "mark as failed" warning work.
- Sourcer pipeline accuracy. Improve "My role = Candidate sourcer" discoverability. He didn’t know this option exists, and the pipeline shows zero phone interviews for them by default (Michael).
- Duplicate Hire profile detection per candidate. Out of pipeline scope. Flag to person details page work (Ella).
- Auto-dismiss for green success banners on the team queue, with accessibility-aware timing. Banners stack visibly when pushing batches (Emily). Worth revisiting the accessibility timing rule with Patty's team.
- Bring back the overview-style homepage with today's events. Q3 concept work already aligns with this. Emily had an immediate positive reaction.

3 Cross-System Data Reach

- Surfaced repeatedly across sessions. Not a single feature, more a roadmap conversation.
- C3 and email. Used for candidate availability confirmation, message-open tracking, and follow-up needs.
- Scheduling page history. Used for decline timestamps (last-24-hour decline detection).
- Interview details page. Used for "interview did not occur" status.
- Candidate profile notes. Used for the recruiter's own notes ("holding on", "great SDE2 when available").
- Debrief tool. Used to surface past debrief content for prep reuse.
- Request notes. Used for Bar Raiser, interviewer, location, and HM. Templated notes (such as GSR's) are easier to start with than free-form notes.
- P0 ships within RPM scope, but the roadmap narrative should explicitly call out cross-system reach. Users have already articulated the gap.

4 SLA and Notification Asks

GSR has SLAs for contact attempts and time-to-schedule. Hire tracks them but does not notify (Ella). NAMER team's 2nd and 3rd contact follow-ups have no notification today (Madhurima). The Q3 SLA-warning roadmap item directly addresses this. AI search becomes significantly more valuable once SLA fields exist as queryable signals.

Adjacent Asks to Surface to PM (not AI search)

- Add HM info to the job rec filter view (Jon).

- Restore referral source visibility (Jon).
- Q3 “My events” feature in Recruiting Hub. Directly enables no-show, decline timestamp, and availability confirmation queries.

Quotable Highlights

- "I don't see a world where if we bring this feature in that I would use filters." — Jon, Recruiter
- "This would be answering off the top of my head kind of all of the issues I would currently see." — Jon, Recruiter
- "It tells me how much of my work is left for the day." — Madhurima, RC Manager (on the result summary line)
- "We just give a freehand text and it generates the results. It's gonna be very, very helpful across for us." — Ubair, RC
- "I will always dream big during these calls, because if I shoot for this, maybe we'll land here." — Kathy, Recruiter
- "I love how simple it is and looks on the screen too." — Ella, RC Lead
- "That is so helpful." — Michael, Sourcer (on discovering "My role = Candidate sourcer")
- "If I could just like magically wish for anything in Hire, it would maybe like highlight the candidates that haven't confirmed their availability." — Camille, RC

Appendix

Test Objectives (Optional Background)

We wanted to learn whether the design successfully communicates seven things:

1. Feature recognition. Do users understand this is an AI-powered natural language search, distinct from the keyword search bar?
2. Input expectations. Do users know what to type and how to phrase queries?
3. System feedback. Does the AI interpretation clearly convey what the system understood?
4. Filter attribution. Can users tell AI-generated filters apart from system filters?
5. Perceived control. Do users know how to modify or undo what the AI did?
6. Role-specific value. Does the feature solve real pain points for each role?
7. Content needs. What do different roles want to search for in their daily work?

Methodology (Optional Background)

- The format was 1-on-1 moderated remote sessions on Chime, 30 to 45 minutes each. The prototype was an interactive HTML build by the designer; data is fake and the UI does not 1:1 match production styling. The protocol combined a walk-through demo with think-aloud and open follow-up questions on workflow needs.
- This was not a quantitative usability test. The goal was directional feedback on whether the design pattern works and what role-specific content needs we should design for.