

Hiring Manager Feedback — Senior Data & AI Consultant

Reviewed by Practice Head (Data & AI Advisory) after evaluating the Draft Shortlist Table, Candidate Summary Cards, and Risk & Gaps Log against the original candidate materials.

Candidate A

Draft Bucket: Advance (8/8 Strong, 0 Gaps)

Decision: Confirm Advance — no adjustment.

Candidate A is the only candidate with a production GenAI/LLM track record on Azure, and it was confirmed hands-on during the screening call. The RAG deployment for the Tier 1 insurance client is directly relevant, and holding both Azure AI Engineer and Azure Data Engineer certs puts them ahead of the pool on certifications without any expiry questions. Consistent C-suite exposure across 2+ years seals it.

Only procedural note: the summary card and draft shortlist credit A with all 8 criteria as Strong. That is fine — the evidence supports it — but ensure the Adaptability rating is traceable. I see 7 years of consulting experience and a record of onboarding onto multiple client projects, which I accept as a proxy for fast ramp-up.

Final bucket: Advance.

Candidate B

Draft Bucket: Advance (8/8 Strong, 0 Gaps)

Decision: Override — move to Hold.

This rating is too generous. The screening notes are unambiguous: "When asked about Azure, said they had 'done some self-study' but no production Azure deployments" and "Acknowledged their background is product-side, not consulting." These are must-have criteria — Cloud Platform Expertise and Industry Experience — not nice-to-haves. Crediting both as Strong is incorrect.

Specific concerns:

- **Cloud Platform Expertise:** B's entire production footprint is AWS. The criterion specifies deep expertise in Azure AI services. Self-study is not a substitute. I am overriding this to No Evidence per the screening record.
- **Industry Experience:** Six years at one SaaS product company is not 5+ years in professional services or technology consulting. The candidate themselves acknowledged the gap. Override to No Evidence.
- **Certifications:** AWS certs are real, but the must-have requires an Azure Data Engineer or Azure AI Engineer. Override to No Evidence — AWS certs do not satisfy this criterion.
- **Stakeholder Management:** "Did not mention any experience presenting to C-suite" per screening notes. No client-facing examples on the resume. Override to No Evidence.

The GenAI/LLM depth is genuine and rare — the AWS Bedrock/LangChain chatbot is impressive work. That plus strong Python are the reasons I want B on Hold rather than Decline. If Candidate A falls through, we could invest in an Azure ramp-up and assign B to a senior engineer on our side for the first client engagement. But B is not ready to lead an Azure AI project independently today.

Final bucket: Hold. Reassess after the next screening round if no other Advance candidate converts to interview.

Candidate C

Draft Bucket: Hold (7/8 Strong, 1 Gap)

Decision: Override — move to Decline.

This is the rating I have the most difficulty with. The draft shortlist shows C with 7 Strong criteria and only 1 gap (GenAI/LLM Deployment). That does not match what I see in the source materials:

- **Programming:** The resume says "Completed an internal Python training programme and started using Python for ad-hoc data cleaning." Screening notes describe skills as "basic — completed an internal training course, not yet used in production." The criterion requires advanced proficiency. Override to No Evidence.
- **Cloud Platform Expertise:** "No cloud certifications and no hands-on Azure or AWS project experience" per screening. Override to No Evidence.
- **Certifications:** None listed on the resume. Override to No Evidence.
- **Project Leadership:** "Could not provide examples of independent project leadership" per screening. Override to No Evidence.
- **Stakeholder Management:** No C-suite interaction referenced. Override to No Evidence.
- **Industry Experience:** Four years in a junior analyst role at a boutique firm does not meet the 5+ year threshold at the required seniority. At best this is Partial, not Strong.

C has no GenAI, no cloud, no Azure certs, basic Python, and no leadership. The screening notes flagged them as "not ready for a senior role" and the candidate themselves described the application as a career-acceleration move. I do not see 7 Strong criteria here — I see maybe 1 or 2 Partial ratings. This candidate should have been placed in Decline based on the evidence.

Final bucket: Decline. Talent Acquisition Lead to keep C on file for a mid-level Data Analyst or Junior Consultant role if one opens.

Candidate D

Draft Bucket: Hold (7/8 Strong, 1 Gap)

Decision: Confirm Hold — no bucket change, but I am flagging two evidence-quality concerns for the next stage.

D has the strongest consulting profile in the pool — nine years in professional services, principal-level, \$1.2M Azure Synapse delivery, established C-suite relationships. No issues there.

The concerns are specific to the AI/GenAI cluster:

1. **Programming:** The draft shows Programming as Strong, which surprised me. D's resume describes data platform work (Synapse, Data Factory, SQL, Power BI) — not Python or AI/ML engineering. I do not see evidence of advanced Python proficiency anywhere. Before the next stage, please confirm what D actually writes code in and at what level. If this is a gap, it belongs as a separate row in the evidence table — D may be a practice lead who directs technical teams rather than a hands-on builder, which changes the fit for this role.
2. **GenAI / LLM Deployment:** The screening notes describe a document summarisation pilot "in early stages and not yet in production" with "high-level" answers when probed. That is Partial — correctly rated. The pilot is promising but we need to push on specifics at the next stage: what architecture, which Azure AI services, and when is production expected? The Azure AI Engineer cert would also help here — follow up on whether that exam has been sat.

I am comfortable leaving D on Hold. If the Programming question resolves positively and the GenAI pilot has real depth, D could move to Advance quickly. If both remain soft, Hold stays.

Final bucket: Hold. Interview preparation to include a technical deep-dive on GenAI delivery and a Python proficiency check.

Summary of Adjustments

Candidate	Draft Bucket	Final Bucket	Reasoning
Candidate A	Advance	Advance	Confirmed — strongest profile, no concerns.
Candidate B	Advance	Hold	AWS-only platform, no consulting experience, no Azure certs. GenAI depth is real but Azure and consulting gaps are must-have failures.
Candidate C	Hold	Decline	Draft over-rated evidence on 6+ criteria. Junior analyst with basic Python, no cloud, no certs, no leadership. Does not meet the bar.
Candidate D	Hold	Hold	Confirmed — but flag Programming evidence quality and push GenAI pilot detail in next stage.

Two candidates (A, D) are recommended for interview preparation (HR04), with Candidate A advancing directly and Candidate D requiring a technical deep-dive. Candidate B is on Hold pending conversion of Advance candidates. Candidate C is declined for this role.

Reviewed by the Practice Head (Data & AI Advisory) on Day 3 of the screening cycle. Cross-referenced against original resumes, screening notes, and must-have criteria from HR01.