

**Role:** Senior Data & AI Consultant **Department:** Data & AI Advisory **Plan ID:** DAA-01 **Debrief Deadline:** Within 48 hours of final interview stage **Submitted Scorecards:** 5 **Panel Members:** 4 unique interviewers across 3 stages

## Scorecard 1: Senior Consultant, Data & AI Advisory

**Stage:** Stage 1 — Technical Screen **Date:** 2026-03-18 **Duration:** 45 minutes

### GenAI / LLM Deployment

Question	Rating	Evidence
Q1: Production deployment walkthrough	Strong	Described deploying a RAG-based document Q&A system for a financial services client. Named Azure OpenAI Service, explained chunking strategy trade-offs (512 vs 1024 tokens), and cited 40% reduction in analyst research time post-launch.
Q2: Deployment failure diagnosis	Strong	Recalled a hallucination issue in production where the model generated plausible but incorrect regulatory citations. Implemented a retrieval validation layer and reduced hallucination rate from 8% to under 1%.
Q3: Production readiness criteria	Strong	Referenced a readiness checklist: automated regression tests, latency monitoring via Azure Monitor, model rollback procedure, and stakeholder demo sign-off before go-live.

### Cloud Platform Expertise (Azure AI)

Question	Rating	Evidence
Q1: Azure AI project walkthrough	Strong	Designed an Azure Machine Learning pipeline for a logistics client. Used Azure Cognitive Services for document OCR and Azure OpenAI for summarisation. Chose Azure over AWS because of client's existing Microsoft EA agreement and data residency in UK South region.
Q2: Performance/cost troubleshooting	Partial	Described an Azure OpenAI token consumption issue that caused unexpected billing spikes, but diagnostic steps were vague — said "we looked at the logs" without specifying which monitoring tools or metrics were used. Resolution was to reduce prompt length, which is a basic fix.

### Programming (Python)

Question	Rating	Evidence
Q1: Complex Python application	Partial	Described building a data pipeline using pandas and FastAPI. Design decisions focused on "keeping it simple" but did not mention testing framework, CI/CD integration, or documentation approach. When probed, said "we did some unit tests but nothing formal."

### Certifications (Azure Data Engineer or Azure AI Engineer)

Question	Rating	Evidence
Q1: Certification influence on practice	Strong	Holds Azure AI Engineer Associate (renewed 2025). Cited the responsible AI module as directly influencing their approach to LLM content filtering in client deployments. Gave a specific example of implementing content safety policies based on certification material.

**Interviewer Summary Notes:** Strong technical depth in GenAI/LLM space — probably the strongest candidate we've seen for this area. Cloud platform knowledge is solid but cost optimisation experience seems limited to basic prompt tuning rather than infrastructure-level optimisation. Python skills are functional but the lack of testing discipline is a gap — may need support from the practice's engineering standards lead during onboarding. Certifications are current and relevant.

---

## Scorecard 2: Practice Head, Data & AI Advisory

---

**Stage:** Stage 2 — Case & Competency **Date:** 2026-03-20 **Duration:** 60 minutes

### Project Leadership (independent client engagements)

Question	Rating	Evidence
Q1: End-to-end engagement leadership	Strong	Led a 6-month AI strategy engagement for a mid-market insurance client. Managed a team of 4, ran weekly steering committees with the client CTO, and delivered on time within a 10% budget variance. Described scope ownership clearly.
Q2: Off-track project recovery	Strong	Identified a scope creep issue 3 weeks into a data platform migration. Renegotiated the statement of work with the client, deferred two low-priority modules, and brought the project back to green status within 2 sprints.
Q3: Knowledge transfer approach	Strong	Described a structured handover framework: documentation repository, 3-session training programme for the client team, and a 30-day post-delivery support window. Client feedback score was 4.5/5.

### Industry Experience (5+ years professional services / tech consulting)

Question	Rating	Evidence
Q1: Client expectations vs. feasibility	Strong	Managed a client who wanted a fully automated underwriting model in 8 weeks. Proposed a phased approach: rules-based MVP in 8 weeks, ML enhancement in phase 2. Client agreed after seeing the risk/benefit analysis.
Q2: Professional services vs. in-house delivery	Strong	Clearly articulated differences: multi-client context, utilisation pressure, knowledge transfer obligation, and the need to build client capability rather than dependency. Gave a detailed example of a 3-month engagement where this dynamic was critical to success.

### Programming (Python)

Question	Rating	Evidence
Q2: Production bug resolution	—	<i>Not asked — ran out of time due to extended case discussion.</i>

### Certifications

Question	Rating	Evidence
Q2: Staying current with Azure	—	<i>Not asked — deferred to Stage 1 coverage.</i>

**Interviewer Summary Notes:** Impressive engagement leadership and client management skills. This is exactly the profile we need to staff our growing enterprise AI advisory work. Strong consulting instincts — understood client dynamics intuitively. I would hire without hesitation.

---

## Scorecard 3: Senior Consultant, Data & AI Advisory

---

**Stage:** Stage 2 — Case & Competency **Date:** 2026-03-20 **Duration:** 60 minutes (same session as Scorecard 2)

## Project Leadership (independent client engagements)

Question	Rating	Evidence
Q1: End-to-end engagement leadership	Partial	Candidate described leading the insurance engagement but when probed on day-to-day activities, the answers focused heavily on technical execution rather than leadership. Steering committee management sounded like it was shared with a delivery manager — unclear who actually owned the client relationship.
Q2: Off-track project recovery	Partial	The scope renegotiation story was compelling at face value, but the candidate attributed the original scope creep entirely to "the client kept adding things" without acknowledging any role in initial scoping. Recovery action was reactive rather than proactive — no evidence of scope management discipline before the issue escalated.
Q3: Knowledge transfer approach	Weak	The structured framework described sounded theoretical. When asked for the specific documentation or training materials produced, the candidate could not name them. Said "the team handled most of the documentation." This suggests delegation rather than ownership of the knowledge transfer deliverables.

## Industry Experience (5+ years professional services / tech consulting)

Question	Rating	Evidence
Q1: Client expectations vs. feasibility	Strong	Same insurance example as the Practice Head noted — the phased approach and risk/benefit framing was genuinely strong. Candidate clearly understood consulting trade-offs when it came to managing client expectations.
Q2: Professional services vs. in-house delivery	Partial	Articulated differences but the example was less specific than it first seemed. When pressed on the "3-month engagement," the timeline kept shifting — initially said 3 months, then said "it might have been closer to 5 months." Raises a minor credibility concern on detail accuracy.

## Programming (Python)

Question	Rating	Evidence
Q2: Production bug resolution	Strong	Described a data pipeline failure where a pandas type coercion silently converted currency values to strings. Implemented type checking at ingestion, added property-based tests with Hypothesis framework, and established a team practice of explicit type annotations. Clear learning and practice change.

## Certifications (Azure Data Engineer or Azure AI Engineer)

Question	Rating	Evidence
Q2: Staying current with Azure	Partial	Mentioned reading Microsoft Learn updates and attending internal lunch-and-learns. Could not name a specific preview feature they had adopted early or a recent platform change that influenced their work.

**Interviewer Summary Notes:** I see a different picture than the Practice Head on leadership capability. The technical skills are genuinely strong, particularly the Python debugging example which showed real depth. But the engagement leadership claims don't hold up under scrutiny — too many instances where the candidate's actual role was unclear or the details shifted under probing. I'd want a reference check specifically on the insurance engagement before I'm comfortable signing off on the leadership assessment.

## Scorecard 4: HR Director, People & Culture

**Stage:** Stage 3 — Leadership & Fit **Date:** 2026-03-22 **Duration:** 45 minutes

## Stakeholder Management (C-suite presentation)

Question	Rating	Evidence
Q1: C-suite presentation	Strong	Presented an AI strategy roadmap to the CEO and CFO of a healthcare client. Framed the value proposition in terms of claims processing cost reduction (projected \$2.1M annual savings) rather than technical architecture. Adapted mid-presentation when the CFO shifted focus to implementation risk — pivoted to a phased rollout slide without losing momentum.
Q2: Stakeholder disagreement	Strong	CTO of the same client wanted to build an in-house ML platform instead of using managed services. Candidate prepared a total cost of ownership comparison and presented it in a 1:1 meeting. CTO agreed to a managed service approach after seeing the 18-month break-even analysis. Evidence-based persuasion without escalation.

## Adaptability (90-day productivity ramp)

Question	Rating	Evidence
Q1: New team productivity ramp		
Q2: Mid-project technology change		

**Interviewer Summary Notes:** Strong communicator with excellent stakeholder management evidence — clear, specific examples with measurable impact. I was called out of the room for approximately 10 minutes during the adaptability section due to an urgent employee relations matter and was unable to complete those questions. I've asked the L&D Manager to cover this area in my absence. Apologies for the incomplete scorecard.

## Scorecard 5: L&D Manager, People Development

**Stage:** Stage 3 — Leadership & Fit **Date:** 2026-03-22 **Duration:** 45 minutes (same session as Scorecard 4)

### Communication Skills

*[Note: This scorecard uses the legacy Q3 2025 ATS competency label. The approved criterion for this interview cycle is "Stakeholder Management (C-suite presentation)."]*

Question	Rating	Evidence
Q1: Communication in team settings	Strong	"Very articulate, good at explaining complex AI concepts in accessible terms. Gave a clear example of running a cross-functional workshop with both technical and business stakeholders."

### Learning Agility

*[Note: This scorecard uses the legacy Q3 2025 ATS competency label. The approved criterion for this interview cycle is "Adaptability (90-day productivity ramp)."]*

Question	Rating	Evidence
Q1: Adapting to new environments	Partial	"Described joining a fintech client mid-project and ramping up on their proprietary data stack within 3 weeks. However, when asked about a structured 30/60/90-day plan, the candidate said they 'don't usually plan it that formally — I just dive in and figure it out as I go.' This approach may not align with Northbridge's structured onboarding programme."
Q2: Responding to methodology changes	Strong	"Gave a strong example of transitioning a team from waterfall to agile delivery during a consulting engagement. Described coaching team members through the change and running retrospectives to identify friction points. Concrete and well-structured answer."

**Interviewer Summary Notes:** Good cultural fit overall. I covered the adaptability questions since the HR Director was pulled away. The candidate's approach to onboarding is a slight concern — the "dive in" attitude works for technical ramp-up but may not align with Northbridge's structured 90-day programme for new consultants. The methodology change example was strong though, suggesting

underlying adaptability is present even if the candidate's default approach is less structured. I used the ATS scorecard template which I now realise still has the old Q3 2025 competency labels — apologies for the inconsistency. Change request CR-2026-0041 to fix the template is still pending.

---

AGASI AiOS — Sample material. For illustrative purposes. All data is fictional.