

HR04 · STEP 4 · FINAL

Interview Prep (Questions & Guides)

## Structured Interview Guides

---

Final interviewer packet with role criteria, question bank, Strong/Partial/Weak rubric, and guidance in one format-consistent package.

### ABOUT THIS DOCUMENT

This is a **sample** deliverable from the **Interview Prep (Questions & Guides)** workflow (HR04) in AGASI AiOS. It shows the expected structure and tone. All names, numbers, and facts are **fictional**.

---

# Structured Interview Guides: Senior Data & AI Consultant

**Role:** Senior Data & AI Consultant **Department:** Data & AI Advisory **Plan ID:** DAA-01 **Interview Format:** Three-stage panel (Technical Screen, Case & Competency, Leadership & Fit) **Must-Have Criteria Source:** Finalized Must-Have Criteria (HR01) **Prepared by:** Talent Acquisition Lead, Talent Acquisition

## Cover Summary

This package contains the complete interview materials for the Senior Data & AI Consultant role at Northbridge Advisory Partners. It compiles the approved question bank, evaluation rubric, and interviewer guidance into a single reference for all panelists.

**Must-Have Criteria assessed in this guide:**

Category	Criterion
Technical Skills	GenAI / LLM Deployment
Technical Skills	Cloud Platform Expertise (Azure AI)
Technical Skills	Programming (Python)
Technical Skills	Certifications (Azure Data Engineer or Azure AI Engineer)
Experience	Industry Experience (5+ years professional services / tech consulting)
Experience	Project Leadership (independent client engagements)
Core Competencies	Stakeholder Management (C-suite presentation)
Core Competencies	Adaptability (90-day productivity ramp)

**Interview structure:**

Stage	Duration	Panel	Criteria Covered
Stage 1: Technical Screen	45 min	Senior Consultant (Data & AI Advisory)	GenAI / LLM Deployment, Cloud Platform Expertise, Programming, Certifications
Stage 2: Case & Competency	60 min	Practice Head + Senior Consultant (Data & AI Advisory)	Project Leadership, Industry Experience, GenAI / LLM Deployment
Stage 3: Leadership & Fit	45 min	HR Director + L&D Manager	Stakeholder Management, Adaptability

**Panel assignment note:** The original ATS configuration assigned the Talent Acquisition Lead to Stage 1. Per the Practice Head's request, Stage 1 is now conducted by a Senior Consultant from the practice team to ensure technical depth. This guide reflects the corrected assignment.

**Scoring note:** The ATS scorecard template still displays legacy prior-cycle competency labels (SQL & Relational Database Proficiency, On-Premise Data Warehouse Management, etc.). Panelists should disregard the legacy labels and score against the criteria and rubric in this guide. Change request CR-2026-0041 to update the ATS template is pending.

---

## Section 1: Approved Question Bank

---

### Technical Skills: GenAI / LLM Deployment

1. Walk me through a GenAI or LLM solution you deployed into a production environment. What was your role, and what were the key technical decisions you made?  
**•What good looks like:** Candidate describes a specific production deployment, names the model architecture or framework, explains trade-offs (cost, latency, accuracy), and states measurable outcomes.
2. Describe a situation where an LLM deployment did not perform as expected in production. How did you diagnose and resolve the issue?  
**•What good looks like:** Candidate identifies the failure mode (e.g., hallucination, latency degradation, data drift), describes a structured debugging approach, and explains the corrective action taken.
3. How do you approach evaluating whether a GenAI solution is ready for production versus still in a prototype stage?  
**•What good looks like:** Candidate references specific readiness criteria (testing coverage, monitoring, rollback plan, stakeholder sign-off) rather than vague statements about "feeling ready."

### Technical Skills: Cloud Platform Expertise (Azure AI)

1. Describe a project where you designed or implemented a solution using Azure AI services. Which services did you use, and why did you choose them over alternatives?  
**•What good looks like:** Candidate names specific Azure AI services (e.g., Azure OpenAI Service, Azure Cognitive Services, Azure Machine Learning), explains the selection rationale, and connects the choice to project requirements.
2. Tell me about a time you had to troubleshoot a performance or cost issue in an Azure AI deployment. What was the root cause, and how did you resolve it?  
**•What good looks like:** Candidate describes a specific Azure environment issue, walks through the diagnostic steps, and quantifies the improvement achieved.

### Technical Skills: Programming (Python)

1. Describe a complex Python application or pipeline you built for a data or AI project. What design decisions did you make to ensure maintainability?  
**•What good looks like:** Candidate discusses code structure, testing approach, dependency management, or documentation practices rather than just describing the output.
2. Tell me about a time your Python code had a significant bug in production. How did you find it and what did you learn?  
**•What good looks like:** Candidate describes the debugging process, the root cause, and a concrete change to their development practice that resulted from the experience.

### Technical Skills: Certifications (Azure Data Engineer or Azure AI Engineer)

1. Which Azure certification do you currently hold, and how has it influenced the way you approach solution design?  
**•What good looks like:** Candidate names a specific active certification, describes how the preparation or knowledge applies to their day-to-day work, and gives a concrete example.
2. How do you stay current with Azure platform changes beyond the certification cycle?  
**•What good looks like:** Candidate cites specific resources (documentation, preview features, community forums, internal knowledge sharing) rather than generic answers.

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

1. Describe a consulting engagement where you had to balance client expectations against technical feasibility. How did you manage the tension?  
**•What good looks like:** Candidate describes the competing priorities, the communication approach used with the client, and the resolution demonstrating consulting judgment rather than just technical execution.
2. What is the most significant difference between delivering AI solutions in a professional services context versus an in-house product team? Give a specific example.  
**•What good looks like:** Candidate identifies structural differences (multi-client context, time-boxed engagements, knowledge transfer requirements) and illustrates with a concrete scenario.

## Experience: Project Leadership (independent client engagements)

1. Tell me about a client engagement you led end-to-end. What did "leading" the engagement look like day-to-day?

•**What good looks like:** Candidate describes scope ownership, client communication cadence, team coordination, and delivery accountability □ not just technical contribution.

2. Describe a situation where a project you were leading went off-track. How did you identify the problem and bring it back on course?

•**What good looks like:** Candidate explains the early warning signs, the corrective actions taken, and the outcome □ demonstrating project management discipline alongside technical skill.

3. How do you handle knowledge transfer to the client at the end of an engagement?

•**What good looks like:** Candidate describes a structured approach (documentation, training sessions, transition plans) rather than ad-hoc handover.

## Core Competencies: Stakeholder Management (C-suite presentation)

1. Describe a time you presented a technical AI solution or strategy recommendation to a C-suite audience. How did you adapt your communication for that audience?

•**What good looks like:** Candidate explains how they translated technical complexity into business impact language, describes the audience's concerns, and notes how the presentation influenced a decision.

2. Tell me about a situation where a senior stakeholder disagreed with your technical recommendation. How did you respond?

•**What good looks like:** Candidate demonstrates active listening, evidence-based persuasion, and willingness to adapt the approach based on valid business concerns □ rather than insisting on technical correctness.

## Core Competencies: Adaptability (90-day productivity ramp)

1. Describe a time you joined a new team or organisation and had to become productive quickly. What did you do in the first 30, 60, and 90 days?

•**What good looks like:** Candidate outlines a deliberate ramp-up approach: relationship building, context gathering, quick wins, and longer-term contribution □ with specific actions at each stage.

2. Tell me about a situation where the technology stack or methodology changed significantly mid-project. How did you adapt?

•**What good looks like:** Candidate describes the change, their learning approach, the time it took to become effective, and how they supported the team through the transition.

## Section 2: Evaluation Rubric

### Technical Skills: GenAI / LLM Deployment

Question	Strong	Partial	Weak
Q1: Production deployment walkthrough	Cites specific model, deployment architecture, trade-offs, and measurable outcomes	Describes a deployment but lacks specifics on trade-offs or outcomes	Cannot describe a production deployment; references only prototypes or coursework
Q2: Deployment failure diagnosis	Identifies failure mode, uses structured debugging, explains corrective action with evidence	Describes an issue but the diagnostic approach is vague or outcome is unclear	Cannot recall a production failure or gives a generic answer with no specifics
Q3: Production readiness criteria	References specific readiness gates (testing, monitoring, rollback, sign-off)	Mentions some criteria but misses key elements like monitoring or rollback	Uses subjective measures ("it felt ready") with no structured framework

### Technical Skills: Cloud Platform Expertise (Azure AI)

Question	Strong	Partial	Weak
Q1: Azure AI project walkthrough	Names specific Azure AI services, explains selection rationale tied to project requirements	Mentions Azure but cannot articulate why specific services were chosen	Cannot name specific Azure AI services or describes non-Azure platforms only
Q2: Azure performance/cost troubleshooting	Describes a specific issue, diagnostic steps, and quantified improvement	Describes an issue but the resolution is vague or unquantified	Cannot recall a troubleshooting scenario in Azure

## Technical Skills: Programming (Python)

Question	Strong	Partial	Weak
Q1: Complex Python application	Discusses design decisions (structure, testing, dependencies) with clear rationale	Describes an application but focuses only on output, not maintainability	Cannot describe a complex Python project or gives a trivial example
Q2: Production bug resolution	Describes the debugging process, root cause, and a concrete practice change	Describes a bug but the learning or practice change is vague	Cannot recall a production bug or describes a minor issue with no learning

## Technical Skills: Certifications

Question	Strong	Partial	Weak
Q1: Certification influence on practice	Names active certification, gives a concrete example of applied knowledge	Names a certification but cannot connect it to daily practice	No active certification, or certification is expired/irrelevant
Q2: Staying current with Azure	Cites specific resources and gives an example of applying new platform knowledge	Mentions staying current but resources are generic ("I read blogs")	No evidence of ongoing learning beyond the certification exam

## Experience: Industry Experience

Question	Strong	Partial	Weak
Q1: Balancing client expectations vs. feasibility	Describes competing priorities, communication approach, and resolution with client	Describes a scenario but the resolution approach is unclear	Cannot provide a consulting-specific example; references only internal projects
Q2: Professional services vs. in-house delivery	Identifies structural differences with a concrete illustrative scenario	Notes some differences but lacks a specific example	Cannot articulate how consulting delivery differs from product teams

## Experience: Project Leadership

Question	Strong	Partial	Weak
Q1: End-to-end engagement leadership	Describes scope ownership, client cadence, team coordination, and delivery accountability	Describes some leadership activities but gaps in end-to-end ownership	Describes technical contribution only; no evidence of engagement leadership
Q2: Off-track project recovery	Explains early warning signs, corrective actions, and measurable outcome	Describes a problem but the recovery approach is reactive or incomplete	Cannot describe a project recovery; blames external factors without personal action
Q3: Knowledge transfer approach	Describes structured approach (docs, training, transition plan)	Mentions handover but the approach is ad-hoc	No evidence of deliberate knowledge transfer

## Core Competencies: Stakeholder Management

Question	Strong	Partial	Weak
Q1: C-suite presentation	Explains business-impact framing, audience adaptation, and decision influence	Describes a presentation but the adaptation for the audience is limited	Cannot provide an example of presenting to senior stakeholders
Q2: Stakeholder disagreement	Demonstrates evidence-based persuasion and willingness to adapt	Describes the disagreement but resolution approach is one-sided	Insisted on technical correctness without engaging with business concerns

## Core Competencies: Adaptability

Question	Strong	Partial	Weak
Q1: New team productivity ramp	Outlines deliberate 30/60/90-day approach with specific actions at each stage	Describes general onboarding but lacks a structured ramp-up plan	No evidence of a deliberate approach; waited to be told what to do
Q2: Mid-project technology change	Describes the change, learning approach, time to effectiveness, and team support	Describes adapting but the learning approach or timeline is vague	Resisted the change or cannot describe how they adapted

## Section 3: Interviewer Guidance

### Introduction

This guide supports a structured behavioral interview format. Each stage targets specific must-have criteria using pre-approved questions. Interviewers should assess candidates against the evaluation rubric provided in Section 2, using the three-level scale (Strong, Partial, Weak) to record evidence-based ratings for each question. Do not deviate from the approved questions during the interview; use the probing follow-ups below to elicit deeper evidence when initial answers are incomplete.

### Recommended Question Sequence

#### Stage 1: Technical Screen (45 minutes)

#	Question Topic	Criterion	Time
1	Production deployment walkthrough	GenAI / LLM Deployment	8 min
2	Deployment failure diagnosis	GenAI / LLM Deployment	7 min
3	Production readiness criteria	GenAI / LLM Deployment	5 min
4	Azure AI project walkthrough	Cloud Platform Expertise	7 min
5	Azure performance/cost troubleshooting	Cloud Platform Expertise	7 min
6	Complex Python application	Programming	5 min
7	Certification influence on practice	Certifications	3 min
□	Buffer / candidate questions	□	3 min

#### Stage 2: Case & Competency (60 minutes)

#	Question Topic	Criterion	Time
1	End-to-end engagement leadership	Project Leadership	10 min
2	Off-track project recovery	Project Leadership	10 min
3	Knowledge transfer approach	Project Leadership	7 min
4	Consulting expectations vs. feasibility	Industry Experience	8 min
5	Professional services vs. in-house delivery	Industry Experience	8 min
6	Production bug resolution	Programming	7 min
7	Staying current with Azure	Certifications	5 min
□	Buffer / candidate questions	□	5 min

#### Stage 3: Leadership & Fit (45 minutes)

#	Question Topic	Criterion	Time
1	C-suite presentation	Stakeholder Management	10 min
2	Stakeholder disagreement	Stakeholder Management	10 min
3	New team productivity ramp	Adaptability	10 min
4	Mid-project technology change	Adaptability	10 min

#	Question Topic	Criterion	Time
□	Buffer / candidate questions	□	5 min

## Probing Follow-Ups

### GenAI / LLM Deployment

- Q1 follow-ups: "What monitoring did you put in place post-deployment?" / "How did you handle model versioning or rollback?"
- Q2 follow-ups: "How long did it take to identify the root cause?" / "What would you do differently if the same failure occurred again?"
- Q3 follow-ups: "Who was involved in the readiness decision?" / "What was the most contentious readiness criterion on that project?"

### Cloud Platform Expertise (Azure AI)

- Q1 follow-ups: "What Azure service did you consider but reject, and why?" / "How did you handle data residency or compliance requirements in Azure?"
- Q2 follow-ups: "What was the financial impact of the issue before resolution?" / "How did you communicate the issue to stakeholders while resolving it?"

### Programming (Python)

- Q1 follow-ups: "How did you approach testing for that application?" / "What would you refactor if you could start over?"
- Q2 follow-ups: "How did you prevent similar bugs going forward?" / "What tooling or process change resulted from that incident?"

### Certifications

- Q1 follow-ups: "Which certification topic has been most directly applicable to your recent work?" / "Has the certification influenced how you mentor others?"
- Q2 follow-ups: "Can you give an example of an Azure feature you adopted early because of your ongoing learning?"

### Industry Experience

- Q1 follow-ups: "How did the client respond to your proposed compromise?" / "What would you do differently in that situation now?"
- Q2 follow-ups: "How does multi-client context affect how you plan your work?" / "What consulting skills do you think are hardest to develop?"

### Project Leadership

- Q1 follow-ups: "How did you handle scope changes from the client mid-engagement?" / "What was your escalation approach when issues arose?"
- Q2 follow-ups: "At what point did you realize the project was off-track?" / "How did you communicate the recovery plan to the client?"
- Q3 follow-ups: "How do you measure whether knowledge transfer was successful?" / "What is the most common failure mode in knowledge transfer you have seen?"

### Stakeholder Management

- Q1 follow-ups: "What question from the C-suite audience caught you off guard, and how did you handle it?" / "How did you prepare for that presentation?"
- Q2 follow-ups: "Did the stakeholder's concern change your technical approach?" / "How did you rebuild alignment after the disagreement?"

### Adaptability

- Q1 follow-ups: "What was the biggest surprise in your first 30 days?" / "How did you identify what to prioritize during ramp-up?"
- Q2 follow-ups: "How did you support team members who were struggling with the change?" / "How long before you felt fully productive with the new stack?"

## Red Flags

- **Vague evidence:** Candidate uses generalities ("I have lots of experience with...") without citing a specific project, outcome, or decision.
- **No production exposure:** Candidate describes only prototypes, demos, or coursework when asked about production deployments.
- **Inability to name Azure services:** Candidate claims Azure expertise but cannot name specific services used or explain selection rationale.
- **Blame shifting:** Candidate attributes project failures entirely to external factors (client, team, tools) without describing their own contribution to the resolution.
- **Leading answer echo:** Candidate parrots the question framing back without adding independent evidence, suggesting rehearsed answers rather than genuine experience.
- **Certification gap:** Candidate cannot name an active, relevant Azure certification or the certification cited has expired.
- **No structured ramp-up approach:** Candidate describes waiting for instructions rather than proactively building context during onboarding, which conflicts with the 90-day productivity target.
- **Resistance to methodology change:** Candidate describes pushing back against technology or process changes without evidence of eventually adapting and contributing.