

Performance Review Narratives

Employee Role: Senior Data & AI Consultant **Department:** Data & AI Advisory — Applied AI Sub-Team **Organisation:** Northbridge Advisory Partners **Review Period:** 14 July 2026 — 14 January 2027 **Prepared By:** Practice Head, Data & AI Advisory (with GenAI drafting assistance) **Date:** 15 January 2027

G1 — Technical Delivery

Rating: Meets

Narrative: Technical delivery goals were met consistently across the review period. The architecture options paper for the GenAI knowledge assistant pilot was delivered three days ahead of the internal deadline and approved with formatting comments only (Evidence Log, Entry 3). The vector database evaluation produced a comparative analysis that directly informed the client's technology selection (Entry 4), and the final implementation report received a 4.5 client satisfaction rating from the Client Technical Lead. One gap was identified: the implementation report initially omitted the fallback strategy for embedding model deprecation, which the client flagged (Entry 8). The missing section was addressed within 48 hours, but catching this in self-review before submission would reduce the risk of client-visible omissions.

Manager Discussion Prompt: What self-review steps could you build into your deliverable process to catch gaps like the fallback strategy before client submission?

G2 — Stakeholder Engagement

Rating: Partially Meets

Narrative: Client relationship-building on the GenAI pilot was effective — the client satisfaction survey returned an average of 4.3/5.0, with the Client Technical Lead and Business Sponsor both rating 4.5 (Evidence Log, Entry 12). However, the steering committee presentation in October ran 8 minutes over the allocated time and a data residency question was deferred to the engagement lead (Entry 5). The VP of Technology's survey comment reinforced this: "Steering presentations could be tighter — we need more time for questions." The consultant's self-assessment acknowledges this gap and references a new practice of preparing pre-read documents. Internally, cross-practice visibility remains limited — the self-assessment notes that relationships outside the Applied AI sub-team have not been systematically developed (Self-Assessment, Section 2).

Manager Discussion Prompt: How are the pre-read documents working in practice, and what other changes to your steering preparation would help you stay within time?

G3 — Business Development

Rating: Below

Narrative: The target of two proposal or scoping contributions per half-year was not met. One contribution was made — the technical methodology section for a retail data strategy proposal — but it required two rounds of revision to meet Northbridge proposal standards, specifically around timeline restructuring and risk mitigation language (Evidence Log, Entry 7). No second proposal contribution occurred during the review period. The self-assessment cites GenAI pilot workload as the primary constraint (Self-Assessment, Section 3). A potential follow-on engagement with the GenAI pilot client was flagged to the Practice Head but no formal follow-up has occurred. The shortfall is partly attributable to the onboarding period, but proactive pipeline engagement was limited.

Manager Discussion Prompt: What would make it easier for you to identify and contribute to proposal opportunities alongside delivery commitments?

G4 — Knowledge Sharing & Team Contribution

Rating: Partially Meets

Narrative: One formal knowledge-sharing session was delivered — the August brown-bag on LLM deployment patterns, which received a 4.2/5.0 average rating from four attendees (Evidence Log, Entry 2). This meets the quarterly target for Q3 but no formal session was delivered in Q4. The planned follow-up session was deprioritised during the GenAI pilot implementation phase (Self-Assessment, Section 4). Informal mentoring of a junior analyst covered code review and client communication over three sessions (Entry 6), though no formal mentoring agreement has been filed. Retrospective contributions were constructive — the engagement onboarding checklist suggestion was accepted and assigned for implementation (Entry 10).

Manager Discussion Prompt: How could you maintain a quarterly knowledge-sharing cadence even during peak delivery periods — would shorter formats or written assets help?

G5 — Professional Development

Rating: Meets

Narrative: The Azure AI Engineer Associate certification (AZ-102) was achieved in December 2026 on the second attempt (Evidence Log, Entry 9). The first attempt in November was unsuccessful, with the consultant identifying responsible AI and governance sections as the gap areas (Self-Assessment, Section 5). The pro-rated learning budget was used appropriately for certification preparation. The long-term development plan discussion with the L&D Manager, originally scheduled for Day 90, was rescheduled to January due to scheduling conflicts and has not yet occurred. The self-assessment expresses interest in a team leadership progression path, which should be discussed during the development planning session.

Manager Discussion Prompt: Now that the certification is complete, what's the next priority for your technical and leadership development — and how can we ensure the L&D Manager discussion happens promptly?

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