

Performance Evidence Log

Employee Role: Senior Data & AI Consultant **Department:** Data & AI Advisory — Applied AI Sub-Team **Organisation:** Northbridge Advisory Partners
Review Period: Current performance review cycle **Compiled By:** Practice Head, Data & AI Advisory **Last Updated:** Current review-cycle close

Log Format

Each entry records a specific observation, deliverable outcome, or stakeholder feedback tied to a date and context. Entries are listed chronologically.

Entry 1

Date: 5 August 2026 **Context:** Predictive maintenance platform engagement (Manufacturing sector) — shadowing phase **Observation:** Completed onboarding onto the engagement faster than expected. Reviewed all existing documentation and attended five squad sessions in the first two weeks. Asked targeted questions about the data pipeline architecture that accelerated their ramp-up. The squad lead (external) noted the new consultant "understood the codebase context quickly."

Entry 2

Date: 22 August 2026 **Context:** Internal — brown-bag session for Applied AI sub-team **Observation:** Delivered a 30-minute presentation on production LLM deployment patterns based on previous experience. Attendance: 14 team members. Feedback collected from four attendees — average rating 4.2/5.0. One comment noted the session "could have included more Northbridge-specific context." Another attendee suggested a follow-up session on cost optimisation.

Entry 3

Date: 12 September 2026 **Context:** GenAI knowledge assistant pilot (Financial Services) — scoping phase **Observation:** Assigned to lead the data architecture workstream for the pilot. Delivered the initial architecture options paper three days ahead of the internal deadline. The Practice Head reviewed and approved with minor comments (formatting only). Client Technical Lead acknowledged receipt but did not provide written feedback.

Entry 4

Date: 3 October 2026 **Context:** GenAI knowledge assistant pilot — implementation phase **Observation:** Led the vector database evaluation and integration workstream. Delivered a comparative analysis of three embedding approaches with clear trade-off documentation. The work directly informed the client's technology selection decision. The engagement lead noted "high-quality technical output" in the weekly status update.

Entry 5

Date: 18 October 2026 **Context:** Client steering committee — GenAI pilot **Observation:** Presented the technical progress update to the client's VP of Technology and two business sponsors. The presentation was well-structured but ran over the allocated time by 8 minutes, reducing time for the Q&A session. The client VP asked a question about data residency that the consultant deferred to the engagement lead. Post-meeting, the engagement lead suggested preparing a pre-read document for future steering sessions to keep presentations within time.

Entry 6

Date: — (date not recorded) **Context:** Internal — mentoring **Observation:** Started informal mentoring of a junior consultant (Analyst grade) in the Applied AI sub-team. Met three times over the first full quarter in role. Focus areas: Python code review practices and client communication. The junior consultant reported the sessions were helpful. No formal mentoring agreement has been filed.

Entry 7

Date: 8 November 2026 **Context:** Proposal support — Retail sector data strategy engagement **Observation:** Contributed to the technical approach section of a proposal for a new retail client. Drafted 4 pages of the methodology section, including the solution architecture and delivery timeline. The proposal was submitted on time. Outcome: pending (as of this log update). The Practice Head noted the draft required two rounds of revision to meet Northbridge proposal standards — primarily restructuring the timeline and adding risk mitigation language.

Entry 8

Date: 22 November 2026 **Context:** GenAI knowledge assistant pilot — client deliverable **Observation:** Submitted the final implementation report for the vector database workstream. Client Technical Lead provided written feedback: "Thorough and well-documented. Minor gap: the report does not address the fallback strategy if the primary embedding model is deprecated." The consultant added the missing section within 48 hours.

Entry 9

Date: 5 December 2026 **Context:** Azure AI platform certification **Observation:** Completed the Azure AI Engineer Associate certification (AZ-102). Certification was achieved on the second attempt — the first attempt (15 November 2026) was unsuccessful. The consultant used 3 days of learning time and the pro-rated learning budget for exam fees and preparation materials.

Entry 10

Date: 12 December 2026 **Context:** Internal — practice retrospective (mid-cycle and year-end) **Observation:** Participated in the bi-annual practice retrospective. Contributed two improvement suggestions: (1) standardise the engagement onboarding checklist across all sub-teams, and (2) create a shared prompt library for common GenAI consulting tasks. The first suggestion was accepted and assigned to the HRIS Administrator for implementation tracking. The second is under review by the Practice Head.

Entry 11

Date: 9 January 2027 **Context:** Quarterly utilisation review — first and second halves of the review period **Observation:** Billable utilisation for the review period: first half — 62%, second half — 71%. Combined average: 66.5%. Below the 70% target. The shortfall in the first half is attributed to the onboarding ramp-up and the 2-week gap between the predictive maintenance engagement wind-down and the GenAI pilot start. Second-half utilisation exceeded target.

Entry 12

Date: 10 January 2027 **Context:** Client satisfaction survey results — GenAI knowledge assistant pilot **Observation:** Client satisfaction survey completed by three client stakeholders. Average rating: 4.3/5.0. Individual ratings: VP of Technology — 4.0, Client Technical Lead — 4.5, Business Sponsor — 4.5. Written comment from VP of Technology: "Strong technical execution. Steering presentations could be tighter — we need more time for questions." No survey was issued for the predictive maintenance engagement (the consultant was in a shadow/support role, not a named team member).

Entry 13

Date: — (exact date not recorded, late in the review period) **Context:** Internal — cross-team interaction **Observation:** The HR Director mentioned during a People & Culture team meeting that the Data & AI Advisory practice had "good examples of knowledge sharing this half." This was not attributed to a specific individual but coincided with the consultant's brown-bag session and retrospective contributions.

Summary Statistics

Metric	Value	Target	Status
Billable utilisation (current review half)	66.5%	70%	Below target
Client satisfaction (GenAI pilot)	4.3 / 5.0	>= 4.0	Meets target
Critical defects	0	0	Meets target
Proposals contributed	1	>= 2 per half-year	Below target
Knowledge-sharing sessions	1 (brown-bag) + retrospective contributions	>= 1 per quarter	Partially meets (1 formal session in 2 quarters)
Certifications	Azure AI Engineer Associate (AZ-102)	Current relevant cert	Meets target