From source files to a coaching briefing

A complete example for coaches, performance analysts and organisers. A fictional Olympic squad brings together athletics, swimming and rowing. The question: what needs follow-up before a nomination review?

What the example actually does

Import the performance CSV and training JSON; select team, season and dates; calculate event-specific progress; add two genuine PDF reports; inspect extracted text; ask a question; retrieve relevant passages; review a cited AI draft; export the briefing with its sources.

Four source files

performance-tests.csv; training-load.json; Squad-Selection-Criteria.pdf; Training-Review-Notes.pdf. All names, measurements, criteria and observations are fictional. The criteria are coach-defined demonstration rules, not official Olympic qualification standards.

Concrete performance findings

2,447 imported measurements; 47 selected for the team and period; 8 athletes across 3 disciplines. Five performance target metrics are met, two need development review and one has insufficient evidence. Four athletes meet their targets and are declared available for staff review. These counts are calculated; they are not predictions or automatic nominations.

A useful document-supported result

The criteria require three comparable tests. The training report records only two tests for Jade. The actual online AI draft identifies the missing third test and proposes arranging it before proceeding with the nomination review. The source passages remain inspectable. This run used 1069 input tokens and 140 output tokens; this is a measurement of this request, not a universal fixed cost.

Online and local are distinct

Operator Web reads these selected files in browser memory. An online AI answer sends the question and explicitly approved selected passages to Vigthoria; those passages may contain personal data. In the separate Windows private workspace, the enabled offline AI processes the question, evidence and answer on the device. Files, approvals and archives do not automatically synchronise between the two.

Windows steps

Install the signed 2.3.2-preview.11 evaluation. Workbench → Private AI → Open local Performance & Coaching. Select both performance files and analyse the squad; add both PDF reports; inspect the text; enter the question; find relevant evidence; analyse with offline AI; check citations; export to an organisation-approved local folder. IT must enable the approved offline AI package first.

Supported formats and bounded processing

Performance session: typed CSV/JSON, 16 files, 1 MiB combined and 10,000 measurements. Documents: text-based PDF, TXT or Markdown, 8 files and 1 MiB combined; 20 pages per PDF; 64,000 extracted bytes per document, 128,000 per session. Scanned PDFs need OCR separately; Word files are not supported by this importer. Retrieval is lexical matching, selecting up to seven source passages; it is not an exhaustive review of every historical record.

Review and deployment

AI summaries are drafts. Citation checks establish that a referenced passage exists; staff must verify that it supports the interpretation. No medical clearance, automated squad decision, prediction, unlimited usage or government certification is claimed. Use managed, encrypted devices and approved export locations; OS software, remote access and sync folders have their own policies. The Windows download is a signed evaluation, not a new Microsoft Store release.