FAQ

MCP for Audit & Assurance

Methodology considerations for using an AI assistant with completed MindBridge analyses

November 5, 2026

VERSION 1.0 | MINDBRIDGE | NOT FIRM METHODOLOGY

IMPORTANT NOTICE

This is MindBridge material. It does not constitute firm compliance, methodology or professional advice. Firms should tailor it to their own methodology, policies and jurisdiction before further distribution.

This FAQ sets out audit methodology considerations for firms and practitioners who use an AI assistant, connected through MindBridge's Model Context Protocol (MCP), to access, query, calculate from or interpret completed MindBridge analyses. It is for audit practitioners and methodology, audit quality and enablement teams.

It does not prescribe firm methodology, documentation, review procedures or quality management policies, which remain matters for the firm and engagement team. For MCP functionality and administration, see the MindBridge MCP overview and Developer Portal.

Three types of MCP output (the audit considerations differ for each)

TABLE 1: TYPES OF MCP OUTPUT

TYPE	EXAMPLES	KEY POINT
Retrieved MindBridge information	Transactions, entries, fields, risk scores, control-point results	Originates in MindBridge
Assistant-generated calculations	Counts, totals, rankings, grouping, filtering, thresholds, trends	Can look factual, but is not necessarily performed or shown by MindBridge
Assistant-generated narrative	Summaries, explanations, themes, observations, suggested risks	Wording, structure and prioritisation are the assistant's

© 2026 MindBridge Analytics Inc. Restricted and confidential.

A. MCP in an audit context

Q1. What does MCP do, and how does it relate to the MindBridge analysis?

MCP is an open standard that lets a supported AI assistant connect to completed MindBridge analyses, so an authorised practitioner can ask questions about them in natural language. The assistant does not perform the analysis: scoring and control-point results are generated within MindBridge, and the **MindBridge user interface (UI) is the source of truth**. MCP is an additional way to interact with the analysis, not an independent version of it.

Q2. Can the assistant change anything in MindBridge?

No. The connection is read-only. The assistant cannot create, modify, delete or run analyses; it can only retrieve information from completed analyses the user is authorised to access. Any filters, calculations or interpretation it applies are not MindBridge results unless they reconcile to the UI.

Q3. Why do prompting and MindBridge terminology matter?

The wording of a request may influence the response. For recurring uses, firm guidance may consider whether prompts specify the engagement, analysis, period, population, filters and output wanted (retrieval, calculation, summary or explanation). Consistent terminology also matters:

- A **transaction** is the broader journal or accounting transaction. An **entry** (or line) is one component of it, so a transaction may contain several debit and credit entries. Counts and scores at the two levels are not interchangeable.
- A MindBridge risk score is not the auditor's assessed risk.
- A control-point trigger is not a conclusion that an item is erroneous, misstated or fraudulent.

Clear prompting reduces ambiguity but does not make a response accurate or complete.

B. Audit evidence and reliance

Q4. When does information obtained through MCP become audit evidence?

MCP is a way of accessing information from a completed MindBridge analysis. Whether any of it becomes audit evidence depends on what is used, what the assistant has done to it, the purpose of the procedure, and the team's assessment of its relevance and reliability. Reliability may also depend on the tools involved: the source data, the MindBridge analysis, the MCP connection and the assistant. The firm's quality management covers whether its technological resources are appropriate, and the engagement team considers their fit for the engagement. For example:

- Locating transactions: whether the correct analysis, period and criteria were used. Evidence more likely comes from follow-up work than the list.

© 2026 MindBridge Analytics Inc. Restricted and confidential.

- Calculating a population: accuracy and completeness of the population and calculation.
- Summarising or explaining results: whether the output is consistent with the UI, including omitted or contradictory information.
- Supporting a significant judgement: whether other evidence supports it, including evaluation of contradictory information.

Q5. Does the assistant simply retrieve results, or can it interpret them?

Both, depending on the request. “Show me the twenty highest-risk journal entries” is retrieval. “Which accounts contain the most high-risk transactions?” involves a calculation. “What themes emerge from the high-risk revenue journals?” produces narrative, where the assistant chooses how to group and prioritise. The underlying results originate from MindBridge; the wording is the assistant's. The more interpretation involved, the more it matters to understand how an observation was derived; significant statements supporting a documented judgement may be verified against the UI.

Q6. Can conclusions be based solely on an assistant-generated summary?

A summary condenses and prioritises, so it may omit lower-ranked items, exceptions or anything the assistant did not treat as responsive. Where it supports a risk assessment or conclusion, considerations include what is omitted and whether significant observations are supported by the UI results, with professional scepticism towards contradictory information. Audit evidence standards look to whether the summary accurately reflects the information on which the judgement rests, not only whether it seems reasonable.

Q7. How can MCP output be traced to MindBridge?

Where a response refers to specific results, the user can look them up in the UI. A factual difference from the UI (populations, scores, control-point results, counts, totals) is a point to investigate, not normal GenAI variability. Possible causes include the analysis or period, transaction versus entry level, filters and assistant calculations.

For a significant conclusion, firms may consider whether a reviewer can trace from the statement to the analysis, the relevant transactions, scores and control points, the criteria applied and the team's reasoning, with query logic and reconciliation for any calculation that does not map to a UI view.

C. MindBridge results and interpretation

Q8. Can the assistant explain why a transaction received a particular score?

It can explain how a named control point works and retrieve a transaction's control-point results. Score-contribution detail through MCP may differ from the UI, which remains the reference. An explanation describes how the analysis reached a result, not whether the transaction is unusual,

© 2026 MindBridge Analytics Inc. Restricted and confidential.

misstated or fraudulent; further investigation may be needed. A zero-weight control point may also trigger without contributing to the risk score, so a trigger does not necessarily explain the score.

Q9. Can the assistant apply firm-specific materiality, thresholds or methodology?

Only if that information is made available to it. Otherwise, “high-risk” may reflect only a MindBridge score range, not the firm's methodology. Where criteria are supplied, considerations include whether the source is approved, current, applicable and correctly applied, which remains for the engagement team.

D. Consistency and reproducibility

Q10. Will the same question always return the same response?

Not necessarily. The MindBridge analysis is reproducible: with source data, configuration, control points, weightings and product version fixed, it contains the same transactions, scores and results. The assistant's response may vary with prompt wording, filters, grouping, row limits or query logic, model updates and, for narrative, prioritisation. A precise request (“all entries scoring above 85”) is more repeatable than “summarise the key risk themes”.

Different assistants may differ in the same ways. Different narrative emphasis may each be reasonable; different factual populations, scores or totals are not alternative interpretations and should be reconciled to the UI (see Q7). Where reproducibility matters, firms may consider keeping the query trace and model version for significant outputs.

E. Methodology, quality and inspection

Q11. Does use of MCP require firms to revisit existing methodology or guidance?

Possibly, though not necessarily through a separate MCP methodology. Existing GenAI guidance may suffice if it covers permitted use cases, the retrieved/calculated/narrative distinction, review, traceability, documentation, and supervision and escalation; if not, firms may consider updating templates, work programmes or review guidance.

Q12. What about data confidentiality and approved assistants?

MCP passes information from completed analyses to the connected AI assistant, so client data and confidentiality obligations apply. Firms may consider which assistants and providers are approved for client data, contractual and jurisdictional restrictions, who may connect an assistant, and data retention or model training. Model versions and updates are likely already covered by general LLM usage requirements. Approval remains a firm decision.

© 2026 MindBridge Analytics Inc. Restricted and confidential.

Q13. How might engagement leadership, supervision and review be affected?

Engagement quality standards place responsibility with the engagement partner, including sufficient resources (which may include technological resources) and tailored direction, supervision and review. Firms may consider approved use cases, competence expectations and escalation criteria.

Review may differ by output type: for retrieval, the analysis, period and criteria; for calculations, the filters, query logic and reconciliation; for interpretive output supporting a significant judgement, the underlying results, omitted or contradictory information and the judgement reflected.

Q14. How can professional judgement remain demonstrably with the team?

Auditing standards require professional judgement and scepticism, and documentation sufficient for an experienced auditor with no previous connection to the engagement to understand the procedures, evidence, results and significant judgements; they do not require retention of every prompt, response or query trace. Depending on significance, firms may consider documenting the purpose of the use, the analysis queried, significant criteria and filters, what was retrieved, calculated or generated, how significant outputs were reviewed, and the rationale for the judgement.

For example, where the assistant identifies themes among high-risk revenue transactions, the file may show how the team considered them against the underlying transactions and assessed risks. If a significant output cannot be re-performed without the prompt or query trace, the narrative alone may be insufficient.

F. Specific uses and practical approaches

Q15. How could a firm use a Skill to support consistent use of MindBridge MCP?

A Skill (or equivalent reusable instruction set, depending on the firm's AI environment) gives the assistant standing instructions. It is optional, supports consistency rather than correctness, and needs testing and ongoing governance. Instructions might cover MindBridge terminology, confirming the analysis and period, labelling output as retrieved, calculated or narrative, and reminding the practitioner that the UI is the source of truth. One approach:

1. Define and test: identify recurring workflows, translate approved firm guidance into reusable instructions with guardrails (avoid audit conclusions), and compare results with the UI.
2. Approve and monitor: set approval and access arrangements, including any internal firm content permitted by information-governance policy, and revisit as tools and guidance change.

© 2026 MindBridge Analytics Inc. Restricted and confidential.