


VERIFIED

A FRAMEWORK FOR TRUSTWORTHY AI ANALYTICS



A METABASE GUIDE

Verified

A practical framework for trustworthy AI analytics

Contents

01	AI moved the analytics bottleneck from creation to curation	4
	Why governance is top of mind for teams now	5
02	Getting started: Six governance questions to ask now	8
	Who gets access to AI analytics?	8
	Where does the data live, and who controls it?	9
	What can AI see?	9
	What happens as usage grows?	9
	Does AI understand your business?	10
	Can you see how AI is being used and what it's doing?	10
03	The five layers of business context	11
04	Treat AI-generated content like code	14
05	The scorecard	17
	Your score	20
06-09	Action plans: Exploring, Building, Governing, Scaling	21
	Exploring (6-9) points	21
	Building (10-13) points	24
	Governing (14-16) points	27
	Scaling (17-18) points	30
10	Data governance is the starting point	33



CHAPTER 01

AI moved the analytics bottleneck from creation to curation.



Actively building
with AI

74%



Say they're scaling

2%

Share of respondents surveyed during Metabase's AI analytics week.

AI made analytics creation easier than ever. Anyone can now ask a question in plain English and get a chart back in seconds.

The harder question that a lot of teams are still figuring out is: where did that data actually go to produce this answer, and whose servers was it processed on? Right behind it is the second question: how do we know this answer is accurate and not an AI hallucination?

Close to half of respondents surveyed during Metabase's AI analytics week said their top requirement for an AI analytics tool was the ability to self-host it or bring their own model and keys, specifically so their data never has to touch a third-party vendor's cloud. In the same survey, 74% said they're actively building with AI, and only 2% called themselves "scaling," meaning they've moved past experimentation to a governed, repeatable practice.

“Creation used to be the bottleneck in analytics, but now it's control, guardrails, and verification.”

Why governance is top of mind for teams now

Where does our data actually go, who can see what, can I trust the number, what happens when usage spikes, does it actually understand our business, and can I see what it's been doing: these are the questions on data leaders' minds now. None of these are solved by a better model. They're solved by the same decisions companies have always made about access, data residency, permissions, and accountability, just now moving at the speed of artificial intelligence.

Data governance is not a new problem, but AI makes the stakes much higher. It moves faster, reaches further, and operates with less friction to stop a bad outcome before it happens. For many companies, this shift turned data sovereignty, or the question of where your data lives, who owns the infrastructure, and who and what gets access, from a once-a-year vendor security review into a question that is required every time someone opens a chat window.

TL;DR

Almost everyone is in motion with AI analytics, but the trust problem is largely still open for most teams. This gap is about data sovereignty, not just who's using the tool. Because adoption is happening faster than governance can keep up, we developed this framework to help you navigate the decisions that determine whether AI analytics can actually be trusted at all — and at scale.



CHAPTER 02

Getting started: Six governance questions to ask now

Begin by understanding what matters most today. Treat these questions as a checklist for any AI-analytics rollout, regardless of which tool you use.

1

Who gets access to AI analytics?

Access to AI analytics shouldn't be a feature toggle you flip for everyone by default. It's a governance decision, the same as any other access to company data, determined by team, by role, or by use case, and revisited as usage grows.

GOOD LOOKS LIKE

Access assigned to defined groups, reviewed on a cadence, not "everyone with a login."

2

Where does the data live, and who controls it?

When AI touches your data, does that data ever leave infrastructure you control? For a lot of teams this used to be a background assumption. Now roughly half of our surveyed respondents named this their top requirement, specifically asking for the ability to self-host or bring their own model and API keys so nothing has to leave their own environment.

GOOD LOOKS LIKE

You have a real choice. You can self-host and bring your own API key to choose which AI model you use when you need to ensure nothing leaves your servers. You can air-gap for a fully offline deployment. Or you can opt for a managed cloud service where that's the right tradeoff.

3

What can AI see?

Separate from where the data lives is who's allowed to see it once AI is asking the questions. Risk goes beyond AI generating wrong answers — it can also be AI surfacing an answer from data someone should not be able to see.

GOOD LOOKS LIKE

AI strictly inherits the existing role-based permissions the person already has, with no parallel permissions or access model to build or maintain.

4

What happens as usage grows?

AI is cheap when five people are using it, and expensive fast once fifty are burning tokens in unison. Cost control isn't optional at scale because it's the difference between a rollout and a runaway bill.

GOOD LOOKS LIKE

Usage limits and alerts set per team, visible before the invoice. Keep token spend in check with limits across the platform, per team and/or per person.

5

Does AI understand your business?

Every company has its own shorthand: what “ARR,” “active,” or “big customer” actually means internally. AI that doesn’t know your definitions doesn’t fail loudly; it gives a confident, wrong answer that sounds right.

GOOD LOOKS LIKE

Your key terms, metrics, and preferred data sources are defined and documented where AI can pull from them, so answers are guided by context to correctness.

6

Can you see how AI is being used and what it’s doing?

Once AI is part of how people get answers, you need the same visibility into that activity that you’d want into any other system producing numbers people act on.

GOOD LOOKS LIKE


Usage analytics and audit logs as a queryable record of who asked what, which tools were used, what data was referenced, what ran, and what it cost.



CHAPTER 03

The five layers of business context

Now that you've set the stage by considering the six questions, it's time to consider if AI understands your business. These are the guardrails that guide AI to correct answers, and it's currently the most underinvested area of AI analytics:



01 **Canonical tables and datasets**

Make the starting point for questions clear for AI by being explicit about which tables it should use, without having to guess which of the three “Customers” tables is the right choice.

02 **Metadata**

Descriptions on your tables and fields, written by someone who understands the data, not inferred by the AI. This is the foundation. Without it, AI is guessing at what a column represents.

03 **Metrics**

Define and save aggregation logic as a metric once, so calculations like ARR are computed the same way wherever they’re queried, by whoever’s asking.

04 **Glossary**

Your company’s internal definitions: what counts as a “big” customer, an “active” user, how revenue is calculated. These are the terms that mean something different in every company, and AI needs them spelled out so it never has to infer or make assumptions.

05 **System prompts**

Standing instructions for how AI should behave: how to answer, what to default to, what to never guess at. This is where you encode judgment calls once so they don’t have to be re-litigated in every conversation.

Get these layers right and reduce your fear of “wrong answers.”

Skip them, and no amount of access control or data sovereignty fixes the fact that AI is confidently answering questions it doesn't understand.




CHAPTER 04

Treat AI-generated content like code

The teams furthest along on the analytics-with-AI journey share one habit: they don't let AI push directly to production. Anything AI builds, from a dashboard to a metric or a report, goes through the same discipline as a code change: a development environment to build and test in, a review step before it ships, and a record of what changed and why.

This matters more as AI gets faster. The easier it is to generate a new dashboard, the easier it is to generate five wrong ones before anyone notices — unless there's a checkpoint between "AI built this" and "a stakeholder is looking at this." It's the essential step that guards against misinformation and endless AI slop proliferating in your organization.

AI CHANGES GO THROUGH REVIEW



“

I’ve always believed in constructive laziness: do the hard work once so you don’t make smart people solve the same problem over and over again. AI should finally make that possible for analytics. The goal isn’t to use AI to crank out ten times more dashboards or automate the same ticket queue we’ve spent twenty years pretending is ‘self-service BI.’

The promise of AI is to let the data team move up a level: define what the numbers mean, curate the data and context, put the right permissions and guardrails in place, make the work inspectable, and then get out of the way.

Because if AI makes analytics creation nearly free, the scarce thing isn’t another chart. It’s having a system where people and agents can get answers without inventing ten different versions of the truth.

The best data teams won’t be the ones answering the most questions, but the ones who’ve built things well enough that they don’t have to.”

**Sameer Al-Sakran**

CEO of Metabase

ABOUT SAMEER

Sameer is the co-founder and CEO of Metabase. He started Metabase in 2014 as a side project while serving as CTO at Expa, building a simple internal tool to help teams get answers from their own data. That instinct became the company's founding bet: open source it, wait years before charging anyone a dime, and grow bottom-up instead of top-down. Over a decade later, Sameer is focused on how AI reshapes analytics and what it takes to build a lasting company in the open.



CHAPTER 05

The scorecard

Now that you understand the tough questions around governance, and have considered the layers required to enable trustworthy AI analytics, evaluate your organization with points (1, 2, or 3) on each question below, then total your points to see where you stand.

QUESTION 1 OF 6

POINTS

Who has access to data for AI analytics?

1 pt

Open to everyone

2 pt

Some teams, no review

3 pt

Assigned by group, reviewed regularly

QUESTION 2 OF 6

POINTS

Data sovereignty and control

- ☐ 1 pt No idea where it runs, default vendor cloud
- ☐ 2 pt Something is in place and assumed fine, but never evaluated
- ☐ 3 pt A real, deliberate choice: self-hosted or bring-your-own-model where it matters

QUESTION 3 OF 6

POINTS

Visibility

- ☐ 1 pt Unsure of what AI can see
- ☐ 2 pt Assume AI follows permissions, but this is unverified
- ☐ 3 pt Strictly inherits existing role-based permissions

QUESTION 4 OF 6

POINTS

Cost control

- ☐ 1 pt No visibility until the bill
- ☐ 2 pt Watched casually, no limits
- ☐ 3 pt Limits and alerts set proactively

QUESTION 5 OF 6

POINTS

Business context

- ☐ 1 pt Not documented, AI guesses
- ☐ 2 pt Known by a few people, might not be written down
- ☐ 3 pt Documented with a semantic layer that’s AI-readable

QUESTION 6 OF 6

POINTS

Auditability

- ☐ 1 pt No record
- ☐ 2 pt Reconstructable manually
- ☐ 3 pt Full queryable audit trail

TOTAL YOUR POINTS

Your score

 / 18

Exploring

6–9 points

Plan → p. 21



Building

10–13 points

Plan → p. 24



Governing

14–16 points

Plan → p. 27



Scaling

17–18 points


Plan → p. 30



Circle your total on the scale, then see your action plan.



CHAPTERS 06-09



ACTION PLAN 06

Exploring

(6-9) points

Exploring	Building	Governing	Scaling
(6-9) points	(10-13) points	(14-16) points	(17-18) points

AI has more access, and your data may be traveling further, than you’ve consciously decided. Start with access and ask where the data actually goes. Do not start by writing an enormous AI governance policy.

Your goal is to reduce the unknowns.

YOUR NEXT 30 DAYS

1

WEEK 1

Map the system.

Document every AI tool currently capable of interacting with company data. Answer the following: Who has access? What data can it reach? Which model/provider processes requests? What information leaves your infrastructure?

2

WEEKS 2-3

Narrow the pilot and test permissions.

Choose one or two groups with a concrete use case. Remove blanket AI access where it isn't needed. Use representative test accounts to confirm that AI cannot retrieve information outside each user's authorized scope.

3

WEEK 4

Establish your first ten definitions and build your glossary.

Document the ten business terms or metrics your pilot group asks about most.

metabase.com/glossary

By the end of the month, you should have:

- ☐ An AI/data-flow inventory
- ☐ Defined pilot groups
- ☐ A permissions test
- ☐ A minimum viable business glossary

Resources

VIDEO



AI privacy

When using our AI service, we don't send the results of your queries to a third-party AI provider, and we don't look at them ourselves. Learn more about what gets sent and to whom.

metabase.com/docs/latest/ai/privacy

VIDEO




Permissions tutorials

Learn how to control who can query specific databases and view or edit shared collections of dashboards and questions by managing group-based access levels.

metabase.com/learn/metabase-basics/administration/permissions/data-permissions


VIDEO



Semantic layer

Wrangle your data, add the context it's missing, and define metrics once so your teammates and AI can query reliably and get the right answer. Get end-to-end visibility: make a change, see what breaks before you ship it.

metabase.com/blog/meet-data-studio-semantic-layer



ACTION PLAN 07

Building

(10–13) points



You’re making deliberate decisions, but one or two layers are still carrying significant risk. The usual trap here is buying more AI capability before fixing the context, observability, or sovereignty underneath it.

YOUR NEXT 30 DAYS

- 1

Find your lowest two scores in the scorecard.

Ignore the dimensions where you’re already strong, and focus on the areas that need improvement.

YOUR LOWEST TWO

1

2
- 2

Turn business context into infrastructure.

Use a semantic layer to create canonical definitions for your highest-impact metrics and add descriptions to your highest-use data sources.

3

Establish cost baselines.

Measure AI usage by team or tenant and decide what should happen when usage exceeds expectations.

4

Start auditing behavior.

Review real questions people ask AI. Those prompts are a roadmap to missing metadata, missing definitions, poor UX, and unmet analytics needs.

5


Assign owners.

Every governance dimension should have a human owner even if the enforcement itself is automated.

By the end of the month, you should have:

- ☐ Canonical definitions for high-value metrics
- ☐ Usage baselines
- ☐ Explicit limits or thresholds
- ☐ Regular AI usage review
- ☐ Named owners

Resources



New in Metabase 61
Build dashboards and questions with MCP

VIDEO

System prompts

System prompts allow administrators to customize instructions, behavior, tone, and domain language for AI.

metabase.com/docs/latest/ai/system-prompts




New in Metabase 61
AI governance and usage analytics

VIDEO

AI usage controls

Granular controls and administrative settings that let you manage who can access built-in AI features, what specific AI tools they can use, and how many tokens or messages they are allowed to spend.

metabase.com/docs/latest/ai/usage-controls




SOURCE NAME	WHERE IT COMES FROM
Metabot	The Metabot chat sidebar inside Metabase.
Documents	Content generation inside Documents.
Suggested Prompts	Background generation of suggested prompts.
Slackbot	Conversations that started in Slack.
SQL	Inline SQL editing in the native editor.
Unknown	A conversation Metabase couldn't classify (distinct from no-source).

DOCS

AI usage auditing

Track and monitor human-robot interactions with AI.

metabase.com/docs/latest/ai/usage-auditing




New in Metabase 59
Data studio

VIDEO

The semantic layer

Wrangle your data, add the context it's missing, and define metrics once so your teammates and AI can query reliably and get the right answer.

metabase.com/docs/latest/data-modeling/data-studio



ACTION PLAN 08

Governing

(14–16) points

Exploring

(6–9) points

Building

(10–13) points

Governing

(14–16) points

Scaling

(17–18) points

Your AI analytics foundation is solid; the remaining gap is more about process than permissions. You know what AI can see and where it runs. Now make sure speed can't bypass the standards you created.

NEXT 30 DAYS

1

Separate development from production for persistent analytics.

AI should be able to experiment without instantly changing what the business relies on.

2

Create a review gate.

Decide which AI-generated assets require human review before production.

3

Add dependency checking.

Before changing shared metrics, transforms, or semantic-layer objects, inspect downstream impact.

4

Version high-value analytics.

Put important production content into a workflow where changes are attributable and reversible.

5

Run one governance fire drill.

Pick a hypothetical problem:

INCIDENT 01

An incorrect AI-generated metric reached a board dashboard

INCIDENT 02

A restricted user retrieved sensitive data

INCIDENT 03

AI spend doubled or tripled

INCIDENT 04


A model/provider became unacceptable overnight

Could you identify what happened and contain it quickly?

By the end of the month, you should have:

- ☐ A dev-to-production path
- ☐ A human review gate
- ☐ Dependency checks
- ☐ Version history for critical analytics
- ☐ An incident playbook

Resources




New in Metabase 57
Remote sync

VIDEO

Git-based workflow

Connect to a git repository so you can push and pull changes to key questions, dashboards, and models just like you manage code.

metabase.com/docs/latest/installation-and-operation/remote-sync



DOCS

Dependency graph

Get a visual map of dependent entities so you can understand what will be affected before making upstream changes.

metabase.com/docs/latest/data-modeling/tools/graph




Metabase AI Agent Skills
metabase-modular-embedding-to-modular-embedding-sdk-upgrade

VIDEO

Agent-driven development

Use a coding agent and the Metabase CLI to create Metabase content, then version that content as YAML files with Remote Sync.

metabase.com/docs/latest/ai/agent-driven-development



ACTION PLAN 09

Scaling

(17–18) points

Exploring

(6–9) points

Building

(10–13) points

Governing

(14–16) points

Scaling

(17–18) points

Your AI analytics are fundamentally sound. Things are well governed, audited, understood. The next challenge is to apply this same discipline as AI moves into more surfaces, more teams, more agents, and potentially your customer-facing product. The danger at this stage is creating one governed AI environment while new ungoverned ones appear around it.

NEXT 90 DAYS

1

Review every new AI interface

against the same six questions above, including:

- In-product AI
- MCP clients
- Coding agents
- Slack or collaboration interfaces
- Embedded customer analytics
- New model providers

2

Measure outcomes, not adoption alone,

because more AI usage is not automatically a success. Track whether AI reduces analyst backlog, increases self-service, speeds decisions, or improves customer experience without increasing incidents or distrust.

3

Turn recurring human review into system design.


If reviewers keep catching the same class of mistake, don't just tell them to keep reviewing it. Ask what semantic definition, permission, system instruction, test, or product constraint could prevent it upstream.

4

Re-score quarterly.

Your architecture, providers, users, agents, and data will change. Governance is not a launch checklist. It's maintenance.

Resources



VIDEO

AI in Metabase overview

AI that meets you where you're at, powered by your preferred model, from Metabot and Slack to MCP, agent-driven development, and the Metabase CLI. Every new interface gets evaluated against the same governance model.

metabase.com/docs/latest/ai/overview



VIDEO

MCP server

If your AI perimeter now includes tools like Codex, Claude or Cursor, understand how agents connect to Metabase and how their access remains scoped to the permissions of the person using them.

metabase.com/docs/latest/ai/mcp



VIDEO

AI auditing and usage analytics

Understand how AI is actually being used in your org, with visibility into which tools are being used, by whom, down to the queries.

metabase.com/docs/latest/ai/usage-auditing



CHAPTER 10

Data governance used to be a checkbox, and now it's the starting point.

AI analytics doesn't require compromising on security, privacy, or control — but it does require answering important questions, and the sooner the better.

If you want [to try Metabase](#), you can bring your own model, including open-weight options. You can host it yourself, and no data has to go to a third-party model provider unless you choose to send it there.

Our isolated workspaces will allow agents to work autonomously against your production warehouse without running amok — sandboxed, with changes reviewed before they touch production.

Because every prompt is exposure: when you send a request to a hosted LLM, it's not just your underlying data that the model sees, it's the prompt itself, logged and retained on that provider's terms, not yours.

The time is now.

Work through the six questions. Fix the areas where your score is lowest. Give each one an owner. Then re-score in a quarter. The goal is to make AI trustworthy for your analytics.

metabase.com/product/data-sovereignty



A METABASE GUIDE

Get started with Metabase

Try Metabase Cloud for free

Try Metabase Cloud →

store.metabase.com/checkout

Take the scorecard online and get your action plan: metabase.com/ai-readiness-hub