

# AI-Powered Analytics is Coming


HOW TO GET YOUR DATA ASSETS READY

You've got an AI-driven analytics initiative. You're exploring agentic AI tools, and the potential is exciting. But before AI can deliver trusted, actionable insights, there's a fundamental question you need to answer:

## **ARE YOUR DATA ASSETS READY?**

# Table of contents

<b>The AI era is here, can you trust the answers?</b>	<a href="#">3</a>
<b>A structured approach for trusted AI analytics</b>	<a href="#">5</a>
Establish and maintain your semantic model	<a href="#">7</a>
Certify your dashboards	<a href="#">10</a>
Connecting your AI agents	<a href="#">12</a>
<b>From us to you: more practical tips</b>	<a href="#">13</a>
<b>Get your data assets AI-ready, and start now</b>	<a href="#">14</a>



# The AI era is here, can you trust the answers?


Generative AI is already changing the way people work, accelerating how teams actually deliver business impact. It's no longer "nice-to-have"—it's a must-have. As to data analytics, the expectation is simple: ask a question, get the right answer.

But making the right answers happen is far from simple.

AI tools don't inherently understand business context. Without clear definitions, they make their best guess, and sometimes, they get it completely wrong.

2.03 PM, 15 Nov 

What's the number of new daily active users added in the past quarter?

 2.03 PM, 15 Nov

To provide an accurate answer I'll need to know which definition of Daily Active User (DAU) you want me to use. I have identified 3 conflicting definitions and similar terms such as: "Active User" and "Engaged User".

## Imagine this:

An AI agent is asked to identify the company's most valuable customers—the premium users.

- ✓ Yes, the “premium” tag exists in the CRM.
- ✓ Yes, AI correctly pulls a list of accounts labeled “premium.”
- ✓ Yes, the data pipeline works as expected.

And yet...

- ✗ No, these aren't actually the company's top customers.

What happened? The “premium” tag was applied six months ago by a marketing campaign and the AI agent stumbled upon it. It didn't “know” that real high-value customers are determined based on **lifetime revenue**, **product engagement**, and **support tier**, not using an outdated campaign label. Worse, the model that defines high-value customers lives in the BI layer together with other exploratory definitions.

LLMs are powerful, but without a trusted source of truth, they can't tell what's outdated and what's accurate. Many think that cleaning the data and removing errors is enough to get AI-driven insights.

But without consistent definitions, even perfectly clean data can still lead to conflicting answers.

The result? Misleading insights, lost trust, and poor decision-making.

This ebook lays out a practical roadmap for ensuring AI-driven insights can be trusted.

Let's dive in →

# A structured approach for trusted AI analytics

To prepare your organization for AI-driven analytics, you need governance. Governed metrics and definitions ensure that dashboards and reports using them are certified as trustworthy.

The key is to create a system where definitions are standardized, data assets such as reports and dashboards are certified, and AI tools can be used with confidence.

The best way to achieve this is by building a **certification program** on top of a **governed semantic model**. This will create a good foundation for trust and consistency.

## But what is a semantic model and a certification program?

### SEMANTIC MODEL

A semantic model is a structured layer that defines and standardizes key business metrics. For GenAI, it powers **text-to-SQL capabilities**. Using the semantic model, AI can accurately answer natural language questions like, “What was our Customer Acquisition Cost (CAC) last quarter?” or, “How many active users did we have last month?”. The semantic model maps business intent to data using standardized definitions, and ensures every answer is accurate, reliable, and aligned with how your organization defines key metrics.

### CERTIFICATION PROGRAM

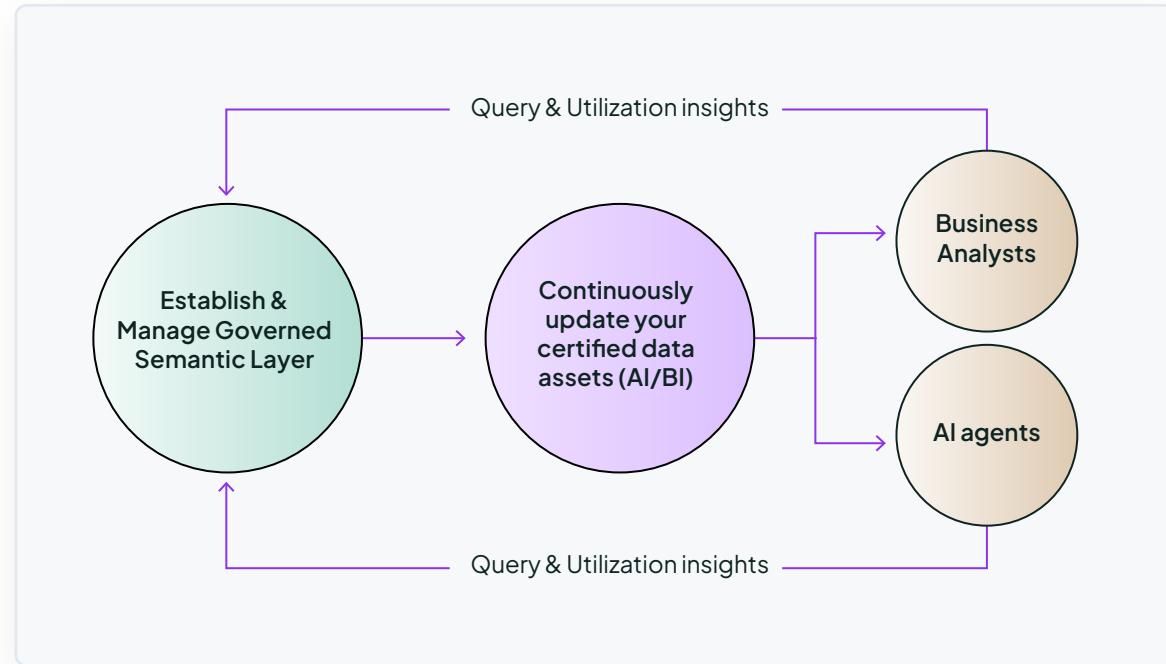
A certification program ensures that data products, such as dashboards or reports, are reliable.

When a dashboard is certified, it means that it meets governance standards and can be shared with confidence. This is what enables **text-to-dashboards/reports**.

Not every dashboard or report can be trusted, even when the organization has a governed semantic model that ensures consistent definitions.

Many dashboards and reports are created for exploratory analysis, quick adhoc reporting, or personal use.

That means that they may not be accurate, up-to-date, or aligned with organizational standards.

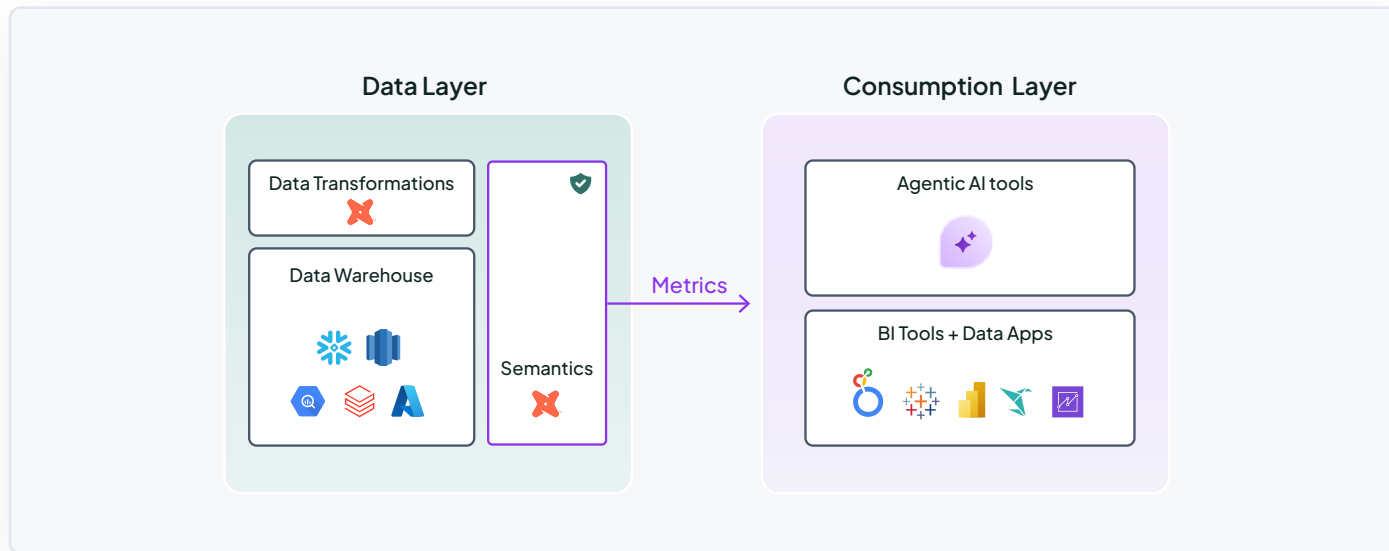


*How to build a certification program on top of a governed semantic model*

A certification program allows AI to surface the most relevant, trusted dashboards, even when multiple versions exist. For example, if a user asks, “Show me our sales performance by region,” AI needs to identify the certified version and present it to the user. The one that’s **up to date, aligned with governed definitions, and free from ad-hoc custom logic**. This ensures business users receive accurate, reliable insights without sorting through inconsistent reports.

Once your metrics and dashboards are governed, your AI agents need to interact with them effectively. AI should pull information only from reliable, certified sources, or at the very least, indicate when a data asset has a low governance score.

**When users learn to trust genAI responses, they’ll stop questioning every number. That’s how you drive real adoption and improve the overall experience of engaging with the data.**



*Implement your semantic model at the data layer to ensure it is shared across all users and applications*

## Establish and maintain your semantic model

There's no shortcut here: having a governed semantic model is critical if you want to implement AI-driven data analytics.

It's the foundation that allows AI tools to interpret business questions accurately and deliver consistent, reliable insights. If text-to-SQL is part of your AI vision, you need a governed semantic model – full stop.

So, how do you get started?

### ① Choose where to build your semantic model

There are many ways to establish a semantic model, but we believe the most effective approach is to **implement it at the data layer**, using tools like dbt semantics. Why? Building at the data layer ensures that your semantic model is shared across all data users and applications. This allows for consistency and alignment across every platform that interacts with your data—whether it's a BI analyst or an AI agent.



















## ② Move key metrics into the semantic model

Not every metric needs to be part of your semantic model. Focus on [metrics that deliver the most value](#) by ensuring consistency and reducing redundancy. Here's how to identify which metrics to prioritize:

- **Highly utilized:** Metrics frequently queried and referenced are likely the most impactful. Governing these first ensures consistency where it matters most.
- **Used across multiple teams:** If a metric is referenced in different reports and dashboards, centralizing it ensures consistency and eliminates discrepancies. Update it once, and it updates everywhere.

- **Critical for long-term decisions:** Prioritize metrics that inform ongoing business performance rather than temporary, project-specific ones.
- **Complex calculations:** Metrics with intricate logic are best defined centrally to ensure accuracy and minimize downstream errors.

Once a key metric is identified, there should be a structured way to integrate it into the semantic model. This could be as simple as opening a ticket when a metric gains traction or defining a workflow for teams to submit new metrics for review.

MOST USED NON-TRIVIAL CUSTOM SQL				<a href="#">View all ↗</a>
 <b>Total Revenue</b> Custom SQL in Sales	 Retention Analysis	 Manny F.	<a href="#">Create Jira ticket</a>	
 <b>Customer Retention Rate</b> Custom SQL in Sales	 User_Behavior	 Chloé G.	6,123 Impressions by 57 users	
 <b>Net Profit Margin</b> Custom SQL in Sales	 Sales Summary	 Chloé G.	5,434 Impressions by 49 users	
 <b>Sales Growth Rate</b> Custom SQL in Sales	 Sales Summary	 Chloé G.	5,234 Impressions by 47 users	
 <b>Lead-to-Customer Ratio</b> Custom SQL in Operational KPIs	 User_Behavior	 Chloé G.	4,831 Impressions by 45 users	
 <b>Profitability Index</b> Custom SQL in Performance	 Customer_Analytics	 Chloé G.	4,031 Impressions by 41 users	

Find key metrics that include complex calculations and create a ticket to add them to the semantic model



### 3 Declutter and resolve conflicts

Identify metrics that are duplicated, unused, or conflicting. Removing what's not needed and resolving discrepancies keeps your semantic environment clean and focused, making it easier for AI to deliver precise answers.

### 4 Monitor your data ecosystem for new definitions

A semantic model isn't static—it evolves with your organization. When a new metric gains traction, evaluate whether it should be added to the semantic model or replace an existing metric. Similarly, deprecate metrics that are no longer relevant.

CONFLICT ANALYSIS: NEW BUSINESS (TABLEAU MEASURE) - 60 DAYS

Name	Parent	Total Impressions	Distinct user impressions	Business domain	Owner
New business	Sales_Policy	4590	99	Sales	Liam Carter
New business	Data	919	40	Sales	Amelia Adams
New business	Custom SQL Query	142	28	Sales	Ella Phillips
new business (1)	Sales_Policy	35	2	Sales	Avery Evans
new business ELIJAH	Custom SQL Query	5	4	Sales	Elijah Walker
new business backup 110324	Custom SQL Query	3	1	Sales	James Anderson
new business Grace (1)	Custom SQL Query	0	0	Sales	Grace Foster
new_business_Lily	Custom SQL Query	0	0	Sales	Lily Morgan
New business	Data	0	0	Sandbox	Noah Mitchell

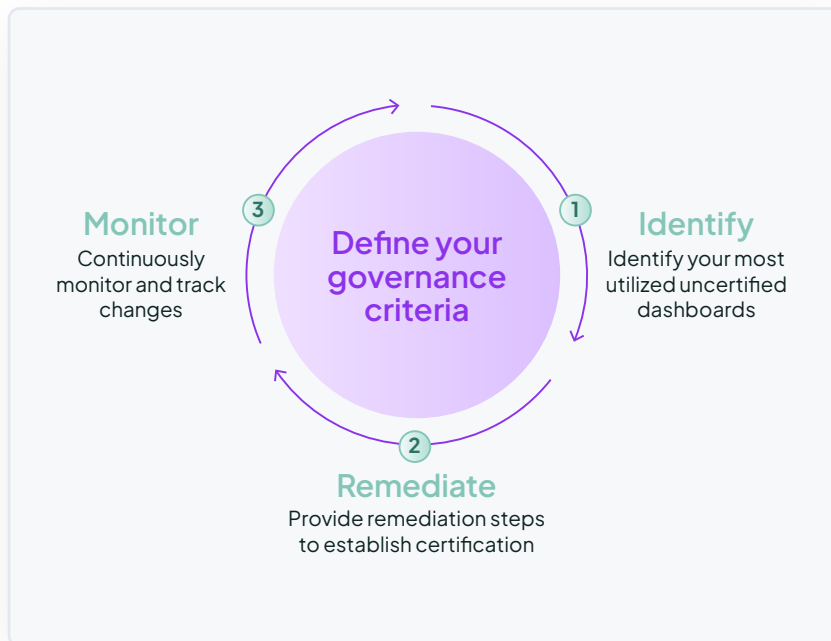
Rows per page: 9 1-9 of 35

euno

Find conflicting measures, trace their data sources and assess their utilization.

## Certify your dashboards

Even with a governed semantic model, AI agents still need to know which dashboards and reports can be trusted. A certification program helps distinguish between reports built on validated, governed data and those created for exploratory or one-off analysis. You can help AI agents navigate dashboards sprawl by implementing a certification program. Here's how to get this done.



### ① Define your governance criteria

Before AI can confidently surface reports, you need clear standards for what makes a dashboard reliable. However, you don't have to enforce strict rules from day one. Start with a simple, flexible framework and refine it over time.

An initial certification framework might include:

- No raw SQL logic embedded in BI tools (avoiding direct, unmanaged SQL queries)
- Built on production-grade tables (such as dbt marts) or other managed layer in the BI tool
- Clearly assigned ownership (so accountability is established for maintenance and updates)

As governance matures, you can begin to introduce stricter criteria, such as:

- Restricting transformations to governed modeling layers
- Enforcing documentation and auditability standards
- Ensuring reports use only definitions from the semantic layer

NON-CERTIFIED HIGHLY USED DASHBOARDS BREAKDOWN View all ↗

<b>Sales Performance</b> Dashboard in Sales	Chloé G.	12 associated views	View recommendations
<b>Marketing Funnel</b> Dashboard in Marketing	Danielle W.	19 associated views	6,847 queries by 65 users
<b>Customer Retention</b> Dashboard in Customers	Ethan B.	21 associated views	5,847 queries by 65 users
<b>Revenue Breakdown</b> Dashboard in Sales	Fiona G.	9 associated views	5,847 queries by 65 users
<b>Product Analytics</b> Dashboard in Analytics	George A.	39 associated views	4,847 queries by 65 users

*Find the most used dashboards that don't meet certification criteria*

## ② Start with the dashboards that matter most

Once your governance criteria are set, you need to verify that dashboards and reports comply. You can do this manually or use tools like [Euno](#) to automate the process.

One of the biggest challenges in certification is knowing where to start. Organizations often have thousands of dashboards and reports, creating sprawl. Many of them are rarely used.

Identifying the most utilized dashboards will let you focus on what matters most. If a highly used dashboard is not certified, provide full visibility into why it doesn't meet governance standards. Notify the owner and offer clear remediation steps. Once you bring this dashboard into compliance, it can be used by AI agents with confidence.

## ③ Continuously monitor and maintain certification

Certification isn't a one-time process.

Data environments are constantly evolving. A dashboard that met governance standards today might become outdated tomorrow due to schema changes, evolving business definitions, or new data sources. Track dashboards continuously and flag those that fall out of compliance to keep AI agents working with reliable data.

MOST QUERIED DASHBOARDS View all ↗

<b>Sales Performance</b> Dashboard in Sales	39 Linked views		Certified	847 queries by 23 users
<b>Customer Segments</b> Dashboard in Sales	14 Linked views		Uncertified	847 queries by 43 users
<b>Churn &amp; Retention</b> Dashboard in Analytics	8 Linked views		Certified	739 queries by 23 users
<b>Influencer Campaign</b> Dashboard in Marketing	3 Linked views		Certified	679 queries by 15 users
<b>Product Lifecycle</b> Dashboard in Inventory	13 Linked views		Uncertified	568 queries by 18 users

*Track your most utilized dashboards and make sure they are certified*

## Connect your AI agents

Your AI-driven data analytics initiative depends on trust and for AI to deliver trustworthy results, it needs to be able to query the governance status of data products and make decisions accordingly. The best way to enable this is by connecting your AI agents to your governance data via an API.

### To answer ad-hoc questions and build predictive models

For queries like, “What was our MRR last year?” or predictive scenarios like, “If we reduce churn by 5%, how will this impact revenue?” AI agents need direct access to governance data, ideally via API, or a natural language interface, so they can pull only governed metrics and definitions. It will ensure every answer is accurate, consistent, and aligned with how the organization defines its key metrics.

### To surface existing dashboards and reports

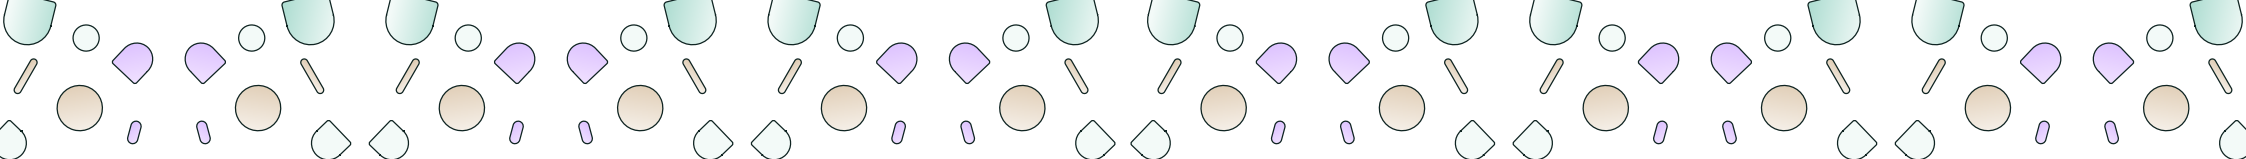
When users ask, “Does this dashboard exist?” or, “Show me our sales performance by region,” AI agents shouldn’t just pull the first report they find. They need to check whether it’s certified and aligned with governance standards. By querying the source of truth, they can confidently surface the most trusted report and avoid irrelevant or outdated versions.

## AI-ready data is simple to do

Getting your data AI-ready doesn’t have to be overwhelming. Start with quick wins: . Invest in your semantic model so AI can reliably interpret business questions and pull the right metrics. Certify your most-used dashboards. Finally, make sure AI agents have access to governance information (whether through an API or a natural language model) so they can confidently surface trusted insights.

For AI to deliver trustworthy results, it needs to be able to query the governance status of data products and make decisions accordingly





# From us to you: more practical tips

Execution matters just as much as strategy. Here are some operational considerations to ensure your AI-powered analytics initiative gains traction and delivers real value.

## **Start where it counts**

Don't try to govern everything at once. Instead, start with a limited set of tables. Choose high-impact areas where trusted data will make the biggest difference such as company-wide KPIs. Then, prioritize domains like revenue reporting, customer analytics, or operational metrics: areas where improved accuracy will immediately benefit decision-making. Once you've built a good foundation and solid processes, expand to your entire data ecosystem.

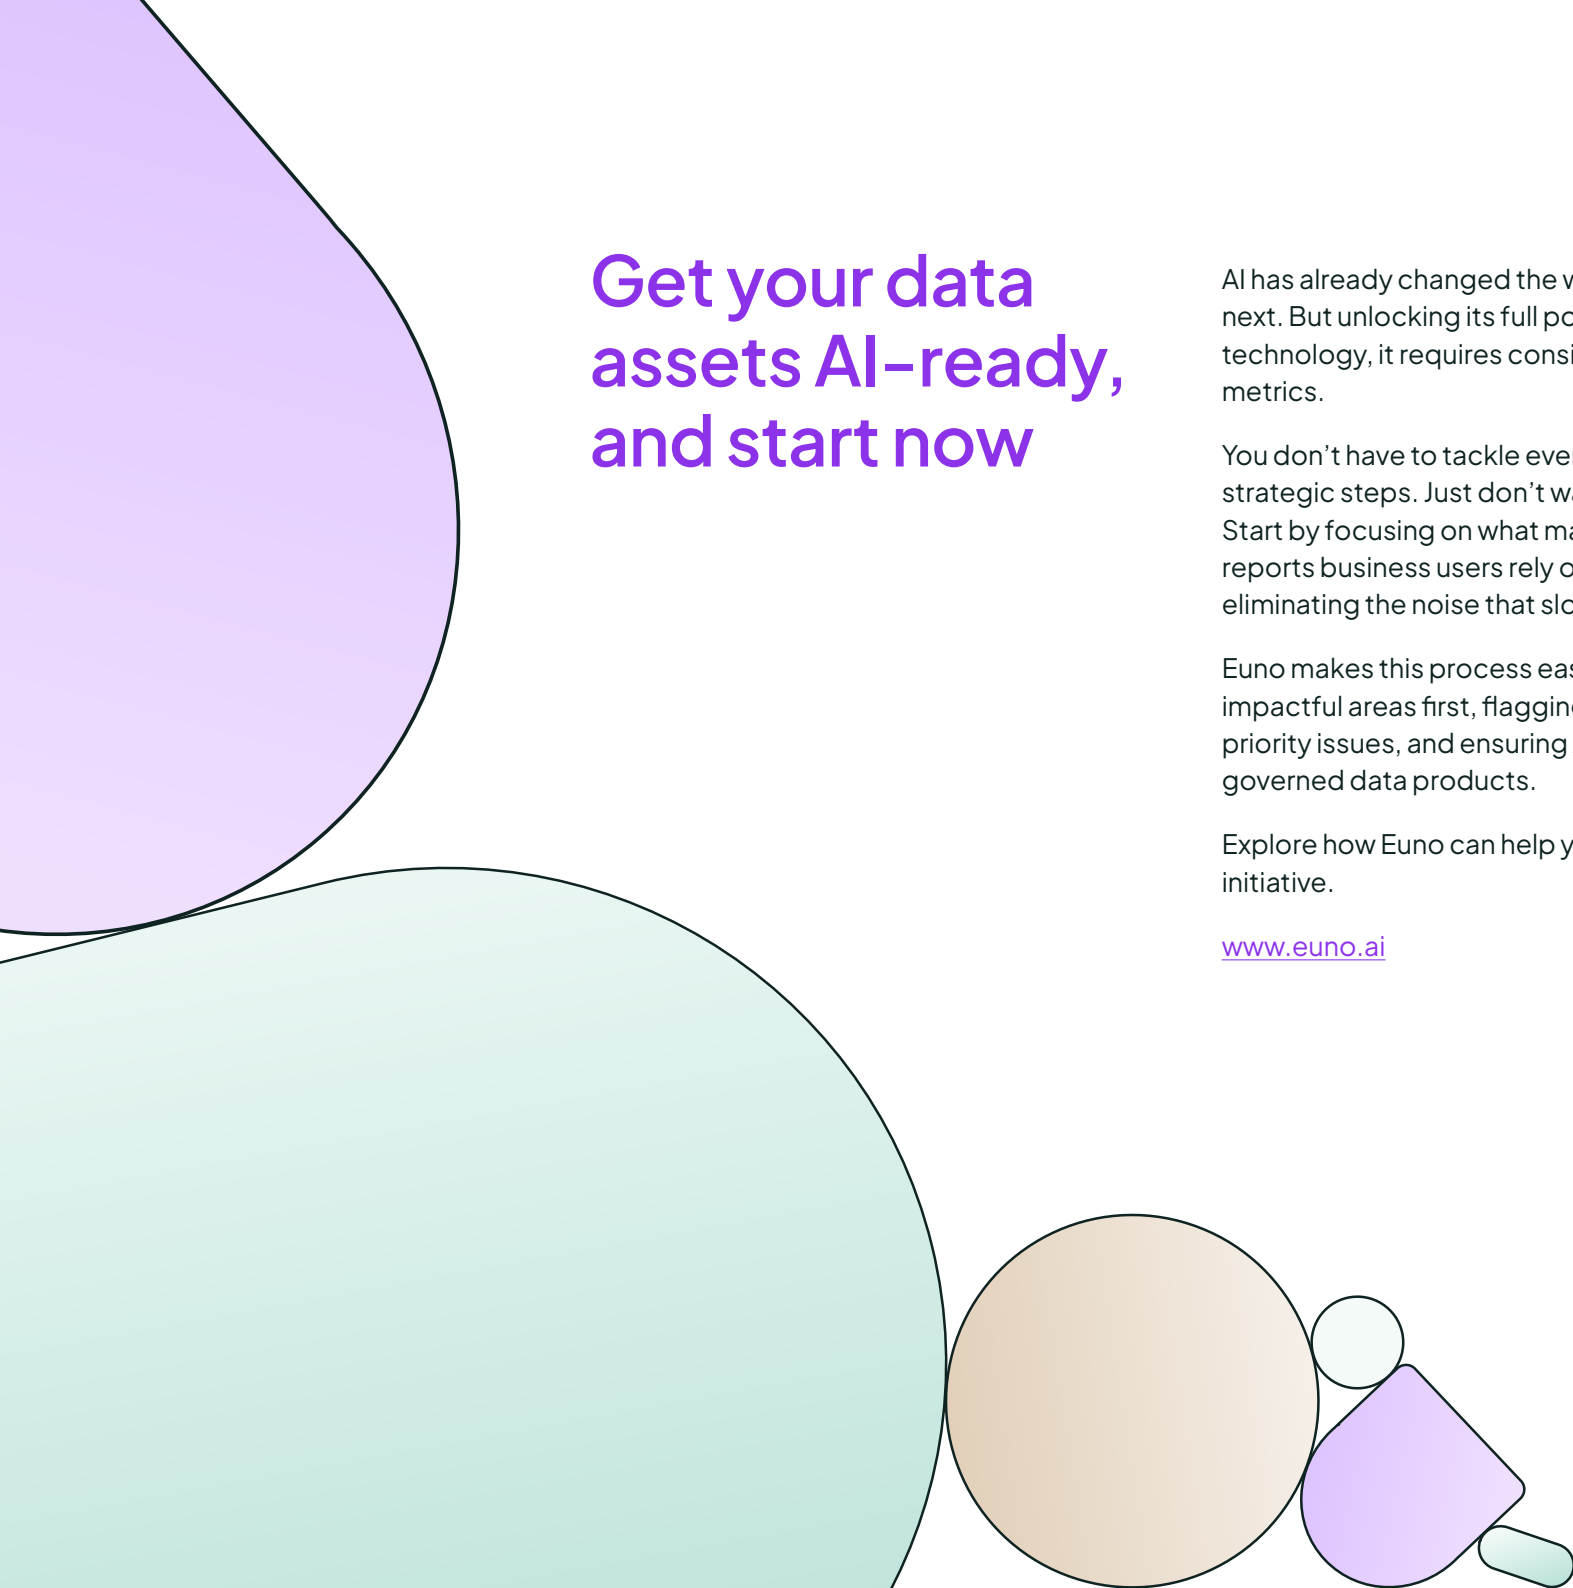
## **Engage your early adopters**

The right first users can make all the difference. Identify data-driven teams who are willing to provide feedback, refine governance processes, and test AI-powered insights. Their input will help iron out issues and validate that certified reports and semantic definitions truly support business needs.

## **Monitor, iterate, and scale**

Governance isn't a one-time setup, it evolves. Continuously track data usage, identify gaps, and refine your certification and semantic model processes. As adoption grows, expand governance to additional domains, ensuring AI continues to operate on high-quality, trusted data.

Follow this and your AI analytics initiative won't just be another experiment, it will deliver real, reliable impact



## Get your data assets AI-ready, and start now

AI has already changed the way we work, and analytics is next. But unlocking its full potential requires more than just technology, it requires consistency on how you define your metrics.

You don't have to tackle everything overnight. Take small, strategic steps. Just don't wait for everything to be perfect. Start by focusing on what matters most: certifying the reports business users rely on, governing key definitions, and eliminating the noise that slows teams down.

Euno makes this process easier by highlighting the most impactful areas first, flagging inconsistencies, surfacing high-priority issues, and ensuring AI agents operate with trusted, governed data products.

Explore how Euno can help you with your AI-powered analytics initiative.

[www.euno.ai](http://www.euno.ai)