
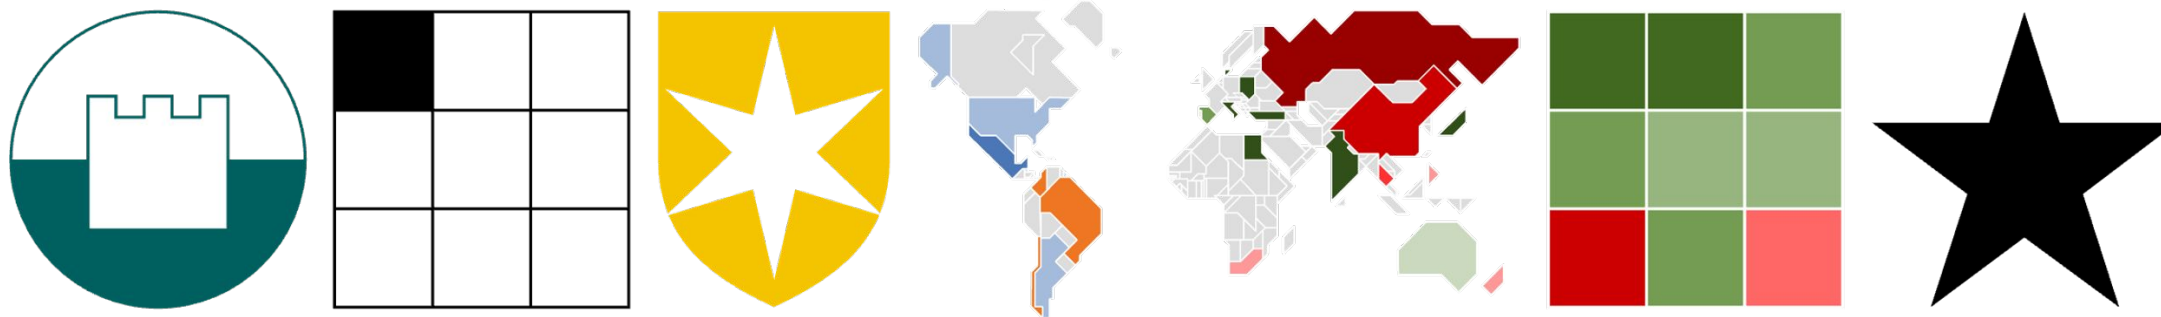




Building Trust in the AI Era of Enterprise Analytics: From Data Integrity to Insight Acceleration

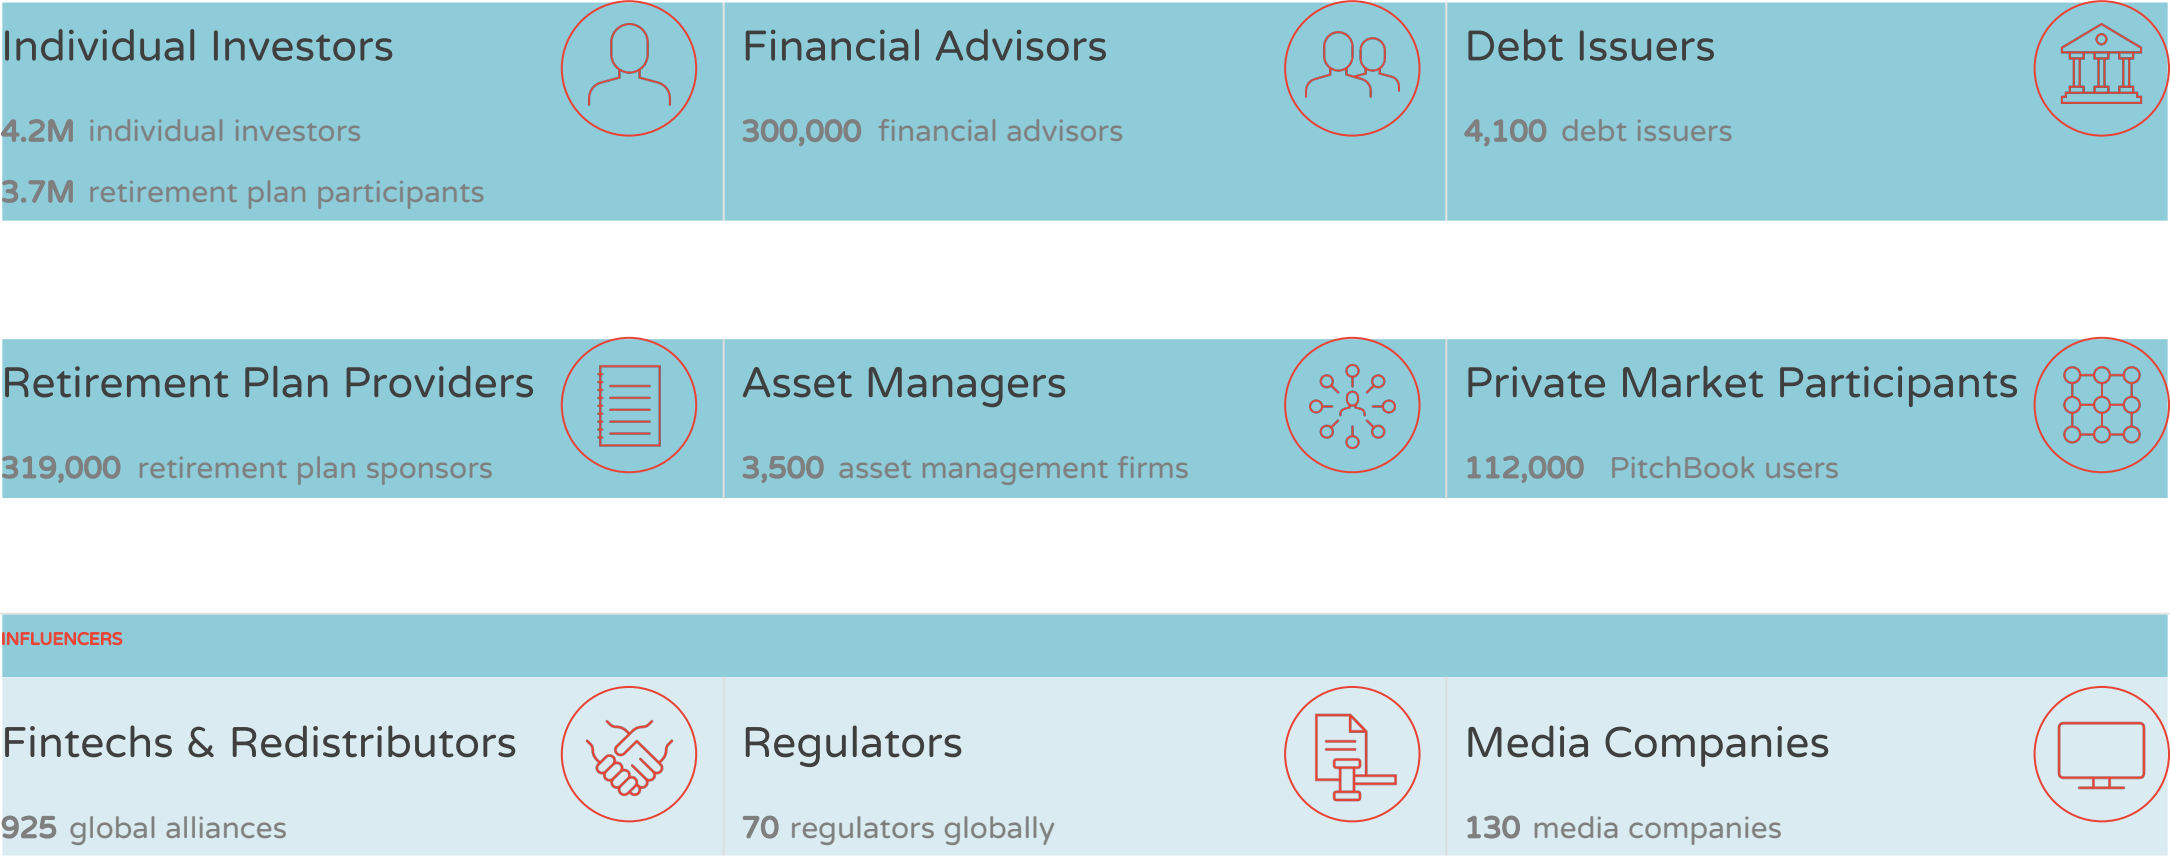
Anusha Dwivedula
Aug 2025





We are the global leader for investor-first insights, tools, and investment strategies, trusted by market participants worldwide. Our system of capabilities, connected through data and insights, is driven by our mission to empower investor success.

Shared language put to work across our broad reach



Data as of Q3 2024

Recent AI Trust Failures

Why Delta Air Lines Is Facing Backlash for AI Pricing

Aug 06, 2025
[Airlines](#) [Delta Air Lines](#) [Opinion](#) [Technology](#) [Travel News](#)



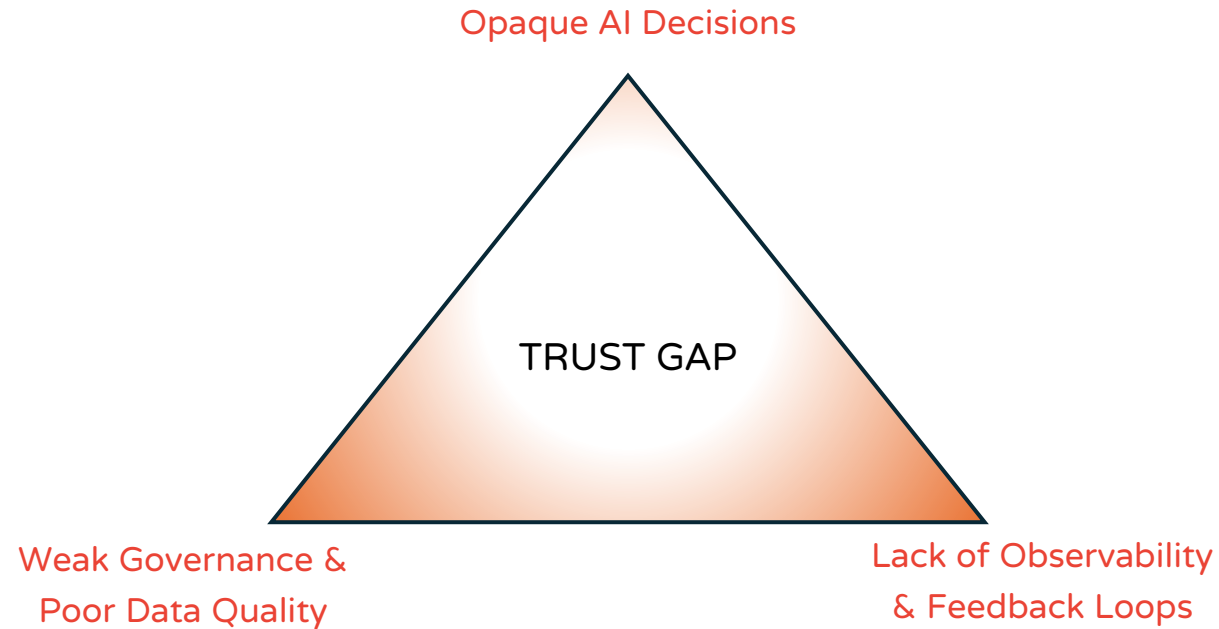
 INDEPENDENT

New AI tool picks up every minor scratch on your car rental – and its freaking people out


James Liddell
Thu 10 July 2025 at 11:05 am GMT-5
3 min read




The Common Thread - The Thing Nobody Talks About Until It's Gone



How do you Build Trust in the AI era of Enterprise Analytics?



AI Driven Quality System



The T.R.U.S.T Pillars



Trust Sandwich Framework




How to Enforce?



Morningstar Total Return Index

- Total return of a fund by tracking the change in net asset value (NAV), reinvesting all income and capital gains distributions, and dividing by the starting NAV.
- Measure the performance of a fund over time, including both price appreciation and the reinvestment of dividends and distributions
- Various applications — Fund and Category Performance, Benchmarking and Investable Products.
- Errors can break client confidence or create regulatory risk.
- Consumers expect to get access to the data within 15 minutes after market




AI-driven quality system



“We needed a system that could be accurate, fast, and trusted—at scale.”

The T.R.U.S.T five pillars of confidence in production systems



The T.R.U.S.T five pillars of confidence in production systems

Transparent — You can trace every decision and every datapoint.

Resilient — It adapts to change and handles uncertainty

Usable — Alerts and insights are clear, actionable, and

aligned with how teams work.
Scalable — It can handle large, complex, fast-changing systems

without slowing down.

Targeted — The system focuses on what matters—no noise,

just meaningful signals.



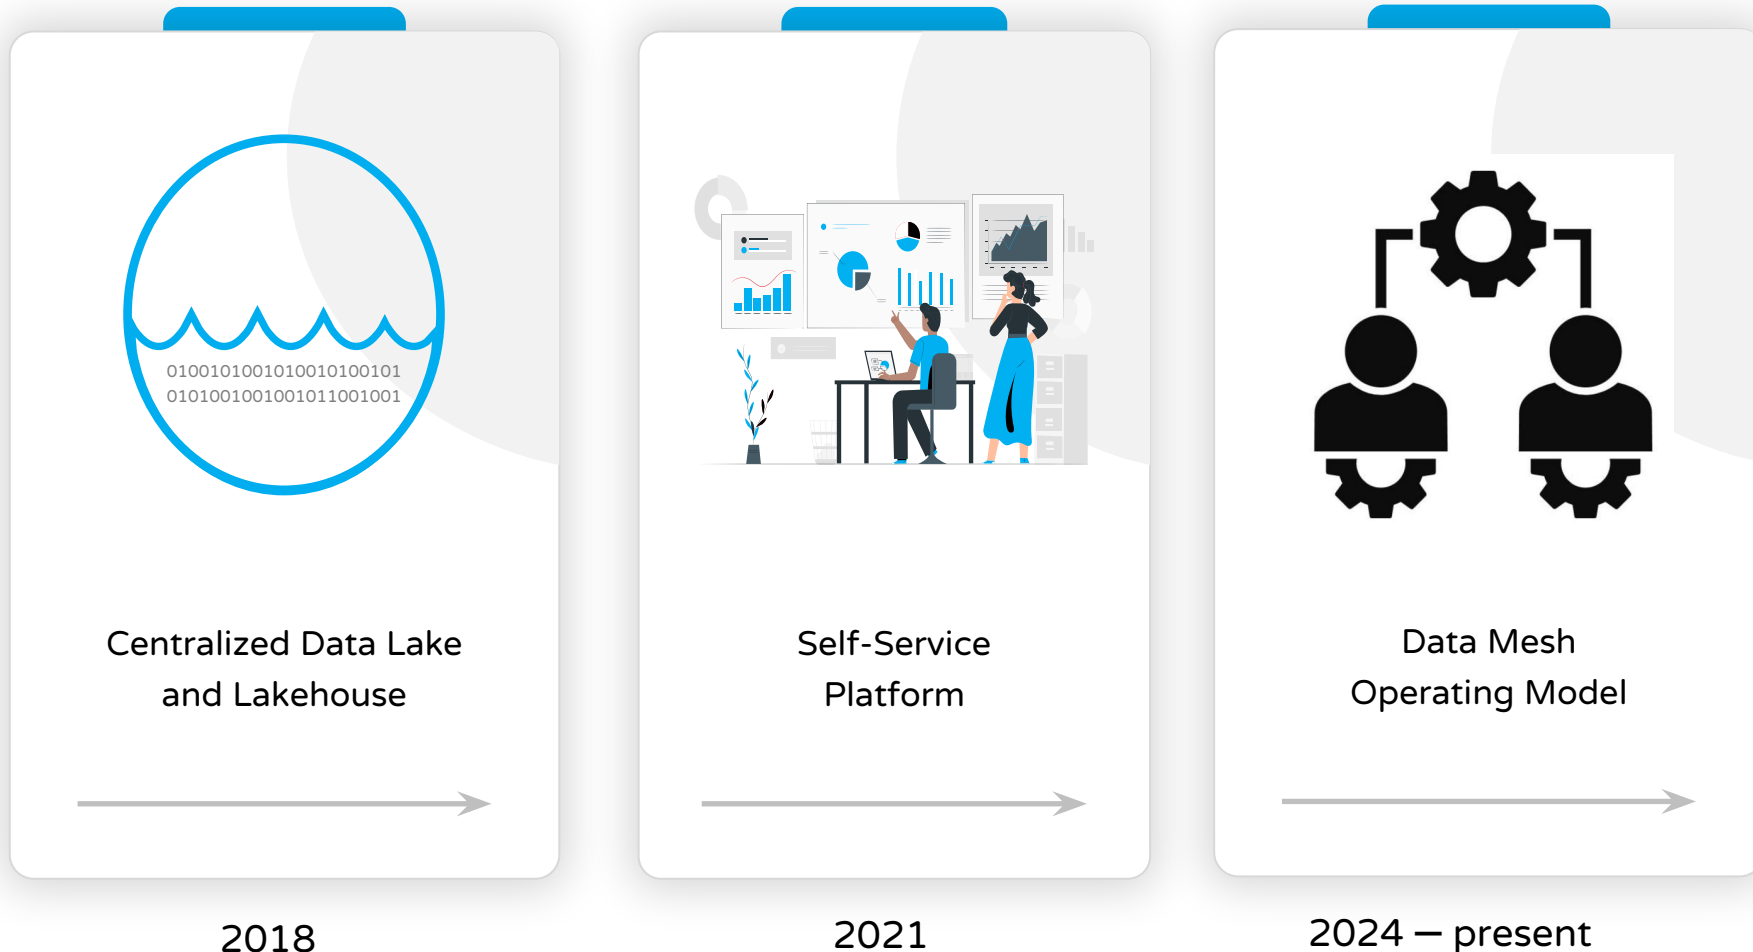
Introducing the TRUST sandwich framework

- Trust isn't built in one place.
- It's a layered system—like a sandwich—where infrastructure, logic, and human validation work together.



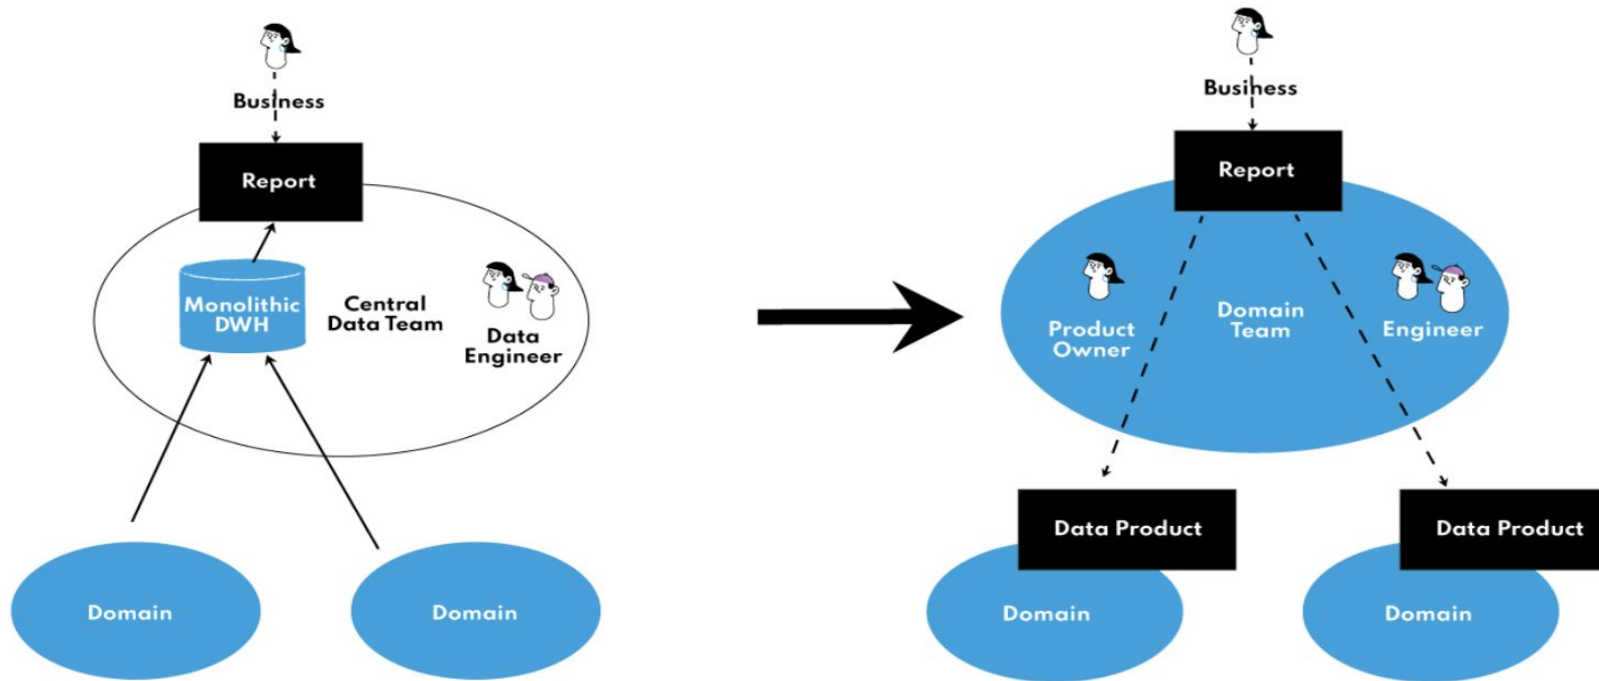
Layer 1: Data Infrastructure and Governance (Transparent and Scalable)

Enterprise Data Platform Journey



Layer 1: Data Infrastructure and Governance (Transparent and Scalable)

De-centralized Sociotechnical Approach

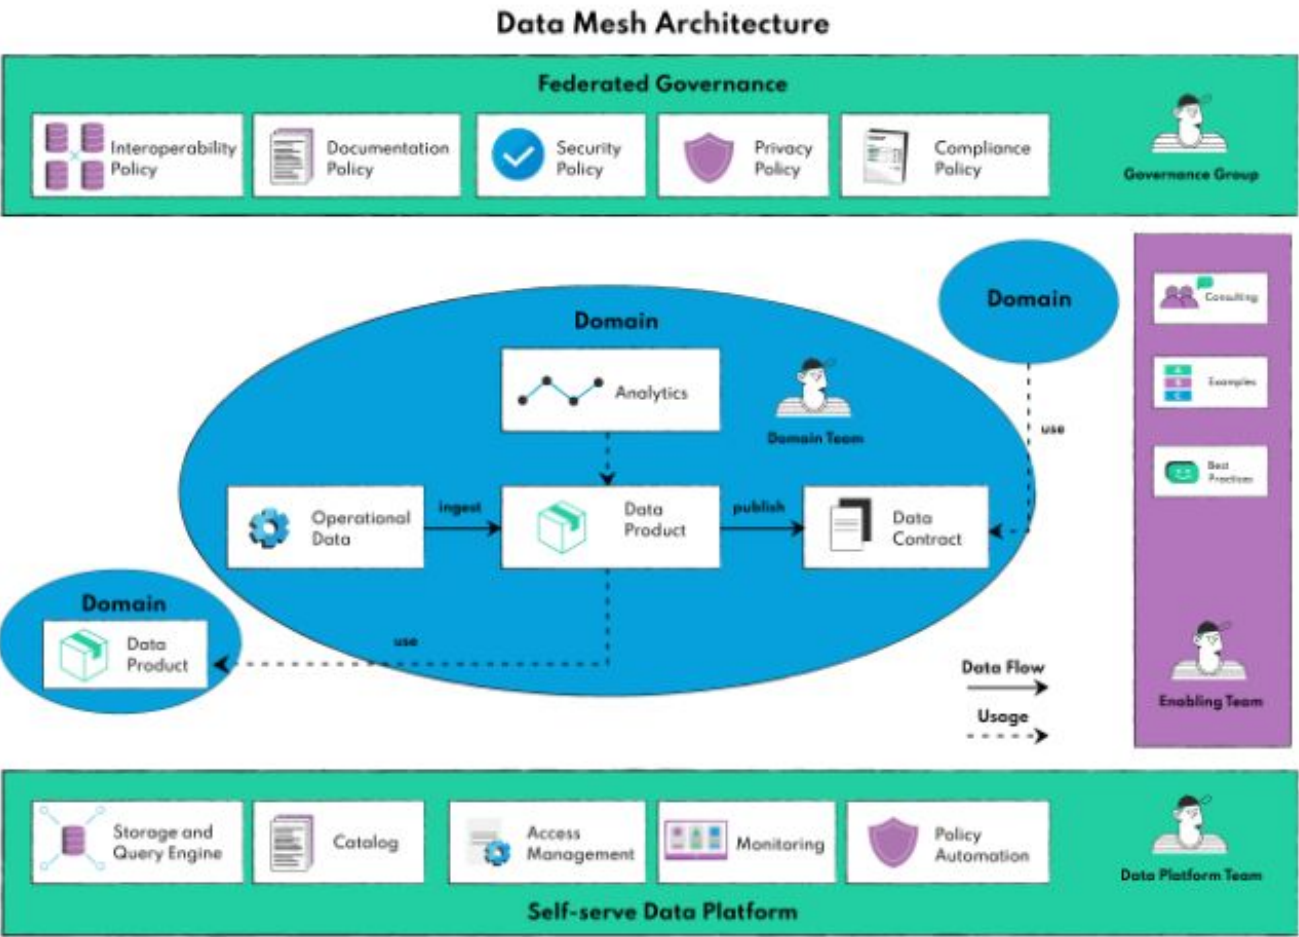


“Data mesh is a decentralized sociotechnical approach to share, access, and manage analytical data in complex and large-scale environments—within or across organizations”

- Dehghani, Zhamak

Layer 1: Data Infrastructure and Governance (Transparent and Scalable)

Data as a Product



Layer 2: AI-driven detection logic and monitoring (Usable and Targeted)

Monitoring Volume of Daily Ingested Data

Algorithm 3 LSTM-Based Data Volume Monitoring

- 1: **Input:** Historical data volumes $V \in \mathbb{R}^n$, external features $X \in \mathbb{R}^{n \times m}$, lag length l
- 2: Normalize V and X using Min-Max scaling
- 3: Prepare lagged sequences of length l for V and X
- 4: Define and compile an LSTM model using Tensor-Flow/Keras
- 5: Train the model on the lagged sequences
- 6: Predict expected volume \hat{v}_t for today
- 7: Compute residual: $r_t = v_t - \hat{v}_t$
- 8: Compute the 95th percentile threshold: $\text{threshold} = \text{Percentile}_{95}(r_t)$
- 9: **if** $|r_t| > \text{threshold}$ **then**
- 10: Flag t as an anomaly
- 11: **end if**
- 12: **Output:** Prediction \hat{v}_t , anomaly status

LSTM is the preferred approach to monitor volume given the need for scalability across different tables*

Algorithm 4 SARIMAX-Based Data Volume Monitoring

- 1: **Input:** Historical data volumes $V \in \mathbb{R}^n$, external features for calendar effect $X \in \mathbb{R}^{n \times m}$, SARIMAX parameters (p, d, q, P, D, Q, s)
- 2: Fit a SARIMAX model: $V_t = \phi(V_{t-1}, \dots, V_{t-p}, X_t, \dots, X_{t-m}) + \epsilon_t$
- 3: Forecast expected volume \hat{v}_t
- 4: Compute residual: $r_t = v_t - \hat{v}_t$
- 5: Compute the 95th percentile threshold: $\text{threshold} = \text{Percentile}_{95}(r_t)$
- 6: **if** $|r_t| > \text{threshold}$ **then**
- 7: Flag t as an anomaly
- 8: **end if**
- 9: **Output:** Prediction \hat{v}_t , anomaly status

SARIMAX is suitable for cases where interpretability and lower computational cost*

*Algorithms cited from IEEE paper we published

Layer 3: Human Context and Expert Validation (Resilient)

Dependency Aware Detection and Customizable Alerts

Algorithm 5 LSTM-Based Anomaly Detection with Expert Opinion

```
1: Input: TRI series  $\Psi \in \mathbb{R}^n$ , lag length  $l$ , threshold  $\theta$ , expert validation set  $E$ 
2: Normalize  $\Psi$  using Min-Max scaling
3: Prepare lagged sequences of length  $l$ 
4: Define and train an LSTM model on historical TRI data
5: for each time step  $t$  in  $1, \dots, n$  do
6:   Predict expected TRI:  $\hat{\Psi}_t = f(\Psi_{t-1}, \dots, \Psi_{t-l})$ 
7:   Compute residual:  $r_t = \Psi_t - \hat{\Psi}_t$ 
8:   if  $|r_t| > \theta$  then
9:     if  $E_t = 1$  (expert confirms valid data) then
10:      Suppress false alarm
11:     else
12:      Flag  $t$  as an anomaly
13:     end if
14:   end if
15: end for
16: Output: List of detected anomalies with expert corrections
```

Expert-guided AI quality control,
unnecessary alerts are reduced,
increasing the system's efficiency*

Algorithm 6 Dependency-Aware Anomaly Detection in TRI

```
1: Input: TRI series  $\Psi \in \mathbb{R}^{n \times m}$ , sector mapping  $\mathcal{I}$ , threshold  $\theta$ , sector-wide threshold  $\tau$ 
2: Normalize  $\Psi$  using Min-Max scaling
3: Group investments into sectors using predefined mappings
4: for each time step  $t$  in  $1, \dots, n$  do
5:   for each sector  $s \in S$  do
6:     Compute sector-wide return  $S_t^s$ 
7:     for each investment  $i \in s$  do
8:       Predict expected TRI:  $\hat{\Psi}^{ti} = f(\Psi_{t-1}^i, \dots, \Psi_{t-l}^i)$ 
9:       Compute residual:  $r_t^i = \Psi_t^i - \hat{\Psi}^{ti}$ 
10:      if  $|r_t^i| > \theta$  and  $|S_t^s - S_{t-1}^s| \leq \tau$  then
11:        Flag  $i$  at time  $t$  as an anomaly
12:      end if
13:    end for
14:  end for
15: end for
16: Output: List of detected anomalies
```

Reducing False Alarms by
Analyzing Data
Dependencies*

Algorithm 7 Customizable Alerting System for TRI Anomalies

```
1: Input: TRI series  $\Psi$ , predicted values  $\hat{\Psi}$ , user-defined severity  $\theta_s$ , persistence level  $p$ , frequency threshold  $\tau$ 
2: Initialize breach counter  $B_w = 0$ 
3: for each time step  $t$  in  $1, \dots, n$  do
4:   Compute residual:  $r_t = \Psi_t - \hat{\Psi}_t$ 
5:   if  $|r_t| > \theta_s$  then
6:     Increment counter  $B_w = B_w + 1$ 
7:     if breach persists for  $p$  consecutive steps then
8:       Trigger persistence-based alert
9:     end if
10:   end if
11:   if  $B_w > \mu_B + \tau$  then
12:     Trigger breach frequency alert
13:   end if
14: end for
15: Output: List of triggered alerts based on user preferences
```

Avoiding Alert Fatigue Through
Customizable Alerting
Mechanisms*

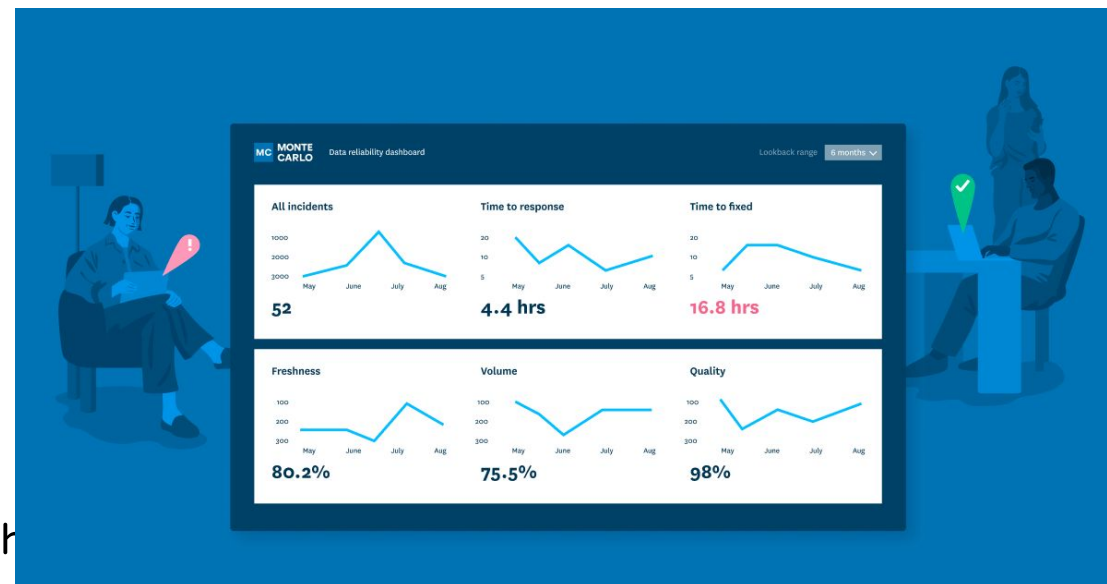
Goal: Alert when it matters!

*Algorithms cited from IEEE paper we
published

Layer 3: Human Context and Expert Validation (Resilient)

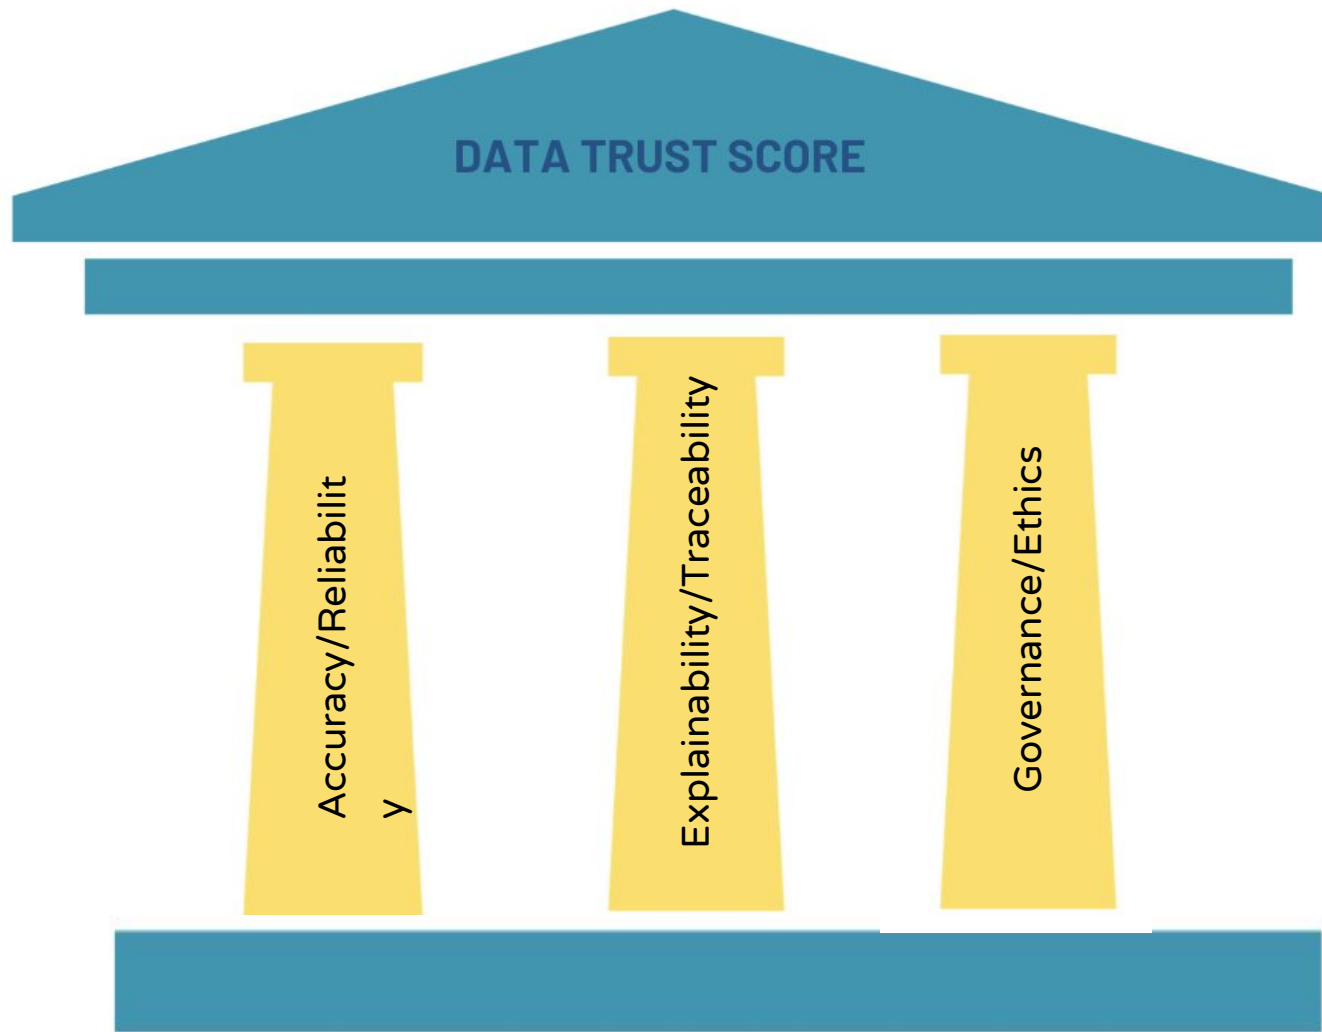
Data Observability

- Experts validate anomalies flagged by the system.
- Confirmed expected fluctuations (e.g., known market patterns)
- Built a feedback loop to refine detection logic over time.
- This gave end-users confidence that the system aligned with reality.



Monte Carlo
Tool

Data Trust Score



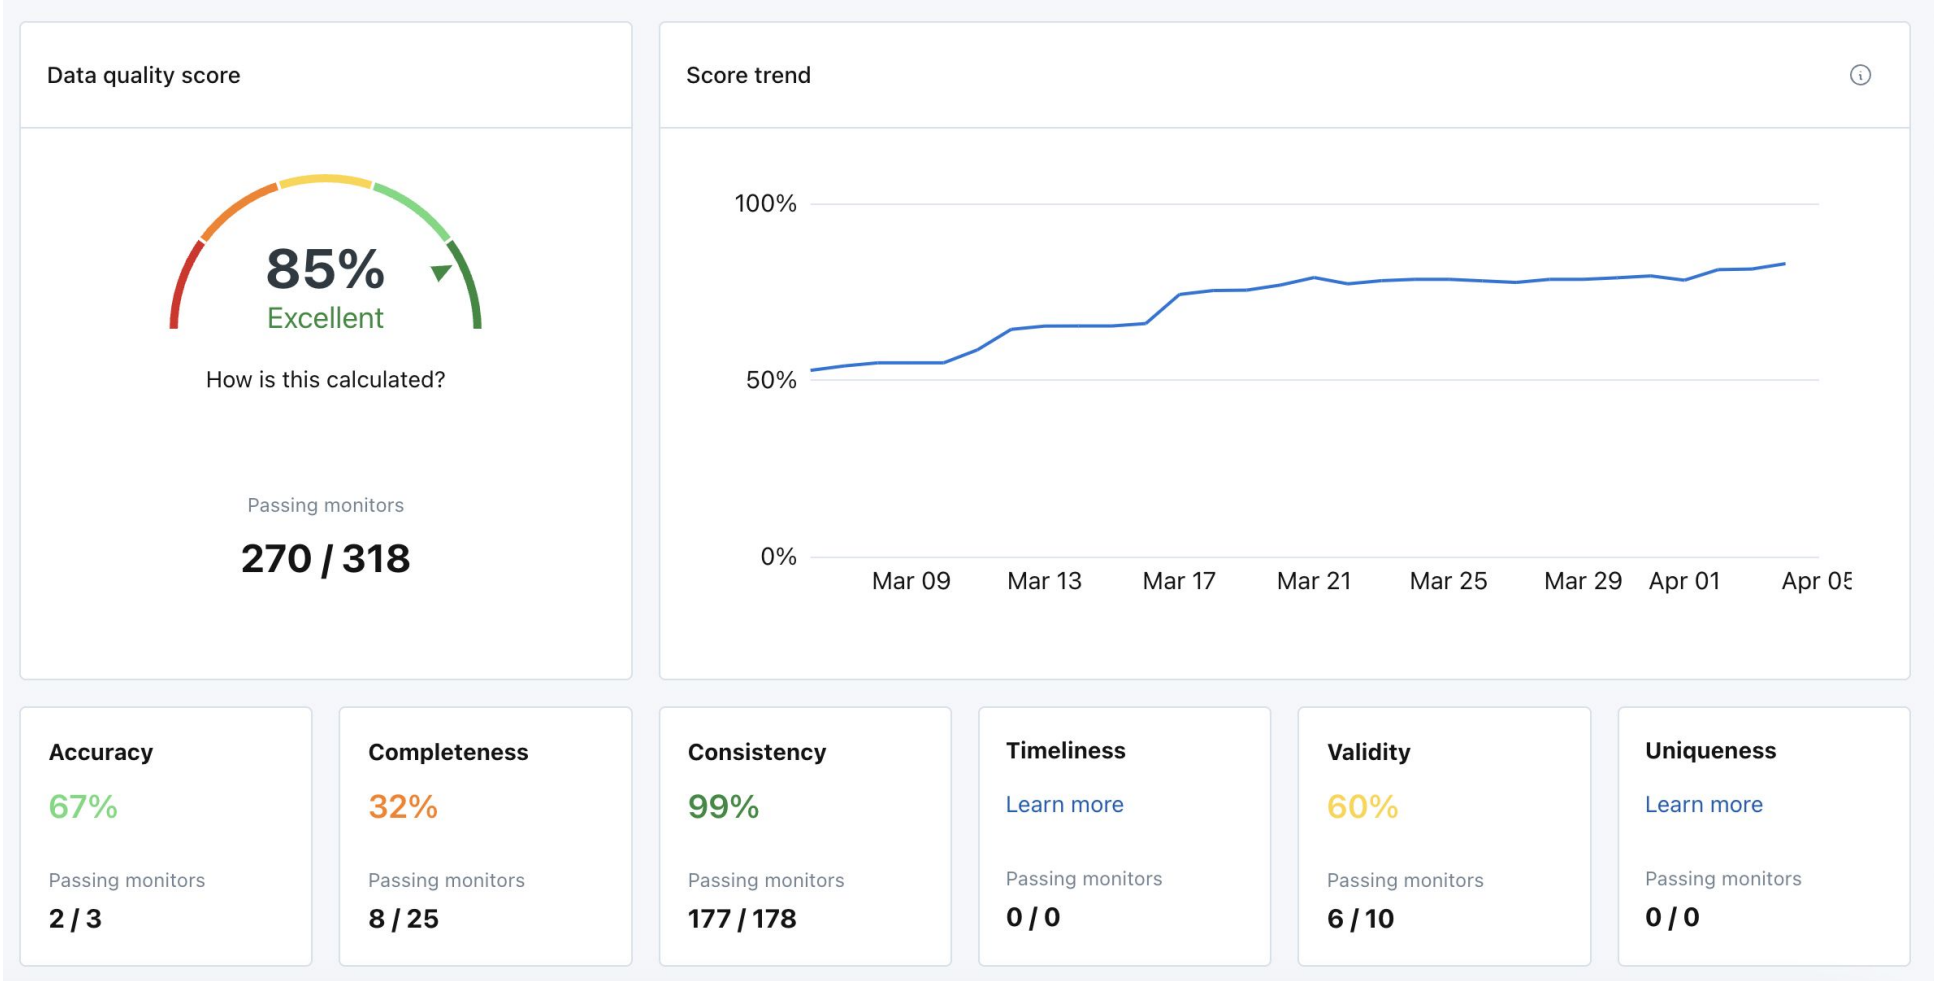
Algorithm 4 Final Data Trust Score Calculation

- 1: Calculate ARS, ETS, and ECGS as described.
 - 2: Aggregate the pillar scores:
 - 3: $DTS = w_1 \times ARS + w_2 \times ETS + w_3 \times ECGS$
-

where w_1 , w_2 , and w_3 are adjustable weights that can be customized based on industry requirements.

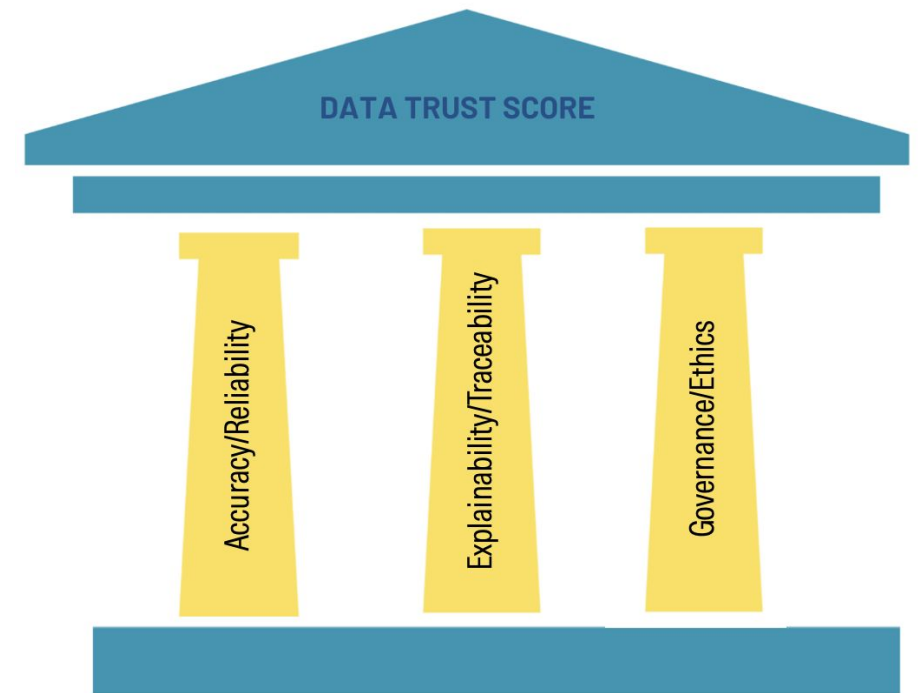
Results

Increased Confidence In Data



“Where is your system leaking trust

today?”
Trust is not a feature—it’s a design decision. And it must be built in layers.



Thank You!

- Establishing Trust in AI-Driven Data Observability and Quality Control: A Framework for Reliable and Scalable Standards
- Trust Erosion: 4 Signals Your Data Strategy Is Breaking Down
(Before AI Fails)
- Personal Website



Anusha Dwivedula

Data Science & Product Leader | Driving AI-Powered Cloud Analytics | Scaling Enterpris...

