

The Modern Data Report

HOW ORGANIZATIONS MANAGE AND SCALE DATA

2024-25

Table Of Contents

About Survey	03
Executive Overview	04
The Cost of Tool Sprawl	05
The Productivity Drain	15
The Data Stack Teams Need	27

What

What is the Modern Data Report?

The Modern Data Report was designed to answer one fundamental question: ***How do today's organizations manage and scale their data?***

The survey reveals how most organizations are struggling to operationalize data. While 75% say they're already working with data products in some shape or form, fewer than half have a clear strategy to execute them. The survey uncovers how fragmentation across tools, misaligned processes, and a lack of governance are limiting their potential.

At its core, the *Modern Data Report* examines how data teams navigate this complexity. It explores key pain points such as inefficient tool ecosystems and governance gaps that hold back useful data-to-insights transformation. The survey is more than a snapshot; it's a call to action for organizations to rethink their approach.

This survey is designed as a yearly roundup of the industry's pulse, tapping into the current challenges, solutions, and opinions of Data Leaders and Practitioners.

2024's report surveyed data professionals from a wide range of industries, including but not limited to financial services, fintech, government, retail, automotive, e-commerce, technology, insurance, manufacturing, banking, healthcare, consulting, education, and edtech.

The respondents hailed from across **48 countries**, the top 5 being the United States, the United Kingdom, Germany, India, and Australia. **232 respondents** participated in the survey with an average of **15+ years of experience** in the data space.

Executive Overview

230+ voices. 15+ years of average experience. 48 countries. One unanimous conclusion: The data stack is crumbling under its own weight. In summary, the survey addressed everything from “What tool should we use?” to “How do we deliver trusted, timely data with minimal friction?”

Time Drain: The Invisible Tax on Data Teams

63% spend more than **20%** of their time on maintenance. **70%** cite business understanding as their biggest time sink. **65%** lose 1/5 of project time just figuring out what data to use. The result? Slower insights. Rework. Missed windows.

The Modern Data Stack is Losing Relevance

The modern data stack is ageing. **55%** set up their infrastructure with multi-tool integrations and now **70%** demand automation to escape the complexity. **85%** say integrating tools across the stack is one of their top three challenges.

Not Aspirations, Real Demands

- **70%** want automation for stack orchestration
- **60%** want integrated data quality
- **50%** need platform-native observability
- **70%** say cataloging beyond data, including infra + services is the next step forward.

The Cost of Tool Sprawl

70% of teams manage **5–10 tools daily**, with over **40%** spending a third of their time simply navigating between them. Integration maintenance is the #1 cost for 40%. Fragmented tools lead to fragmented truths, creating confusion across the data ecosystem.

What the Data Stack Needs Now

The demands are unanimous and clear-cut.

- Ability to abstract tool integration
- Centralize governance while enabling decentralized execution
- Native metadata and observability
- Semantic search, natural language querying, and business-aligned data exploration

Model-First, Product-Centric Future

65% say the twin engines of acceleration are **data models and data products**. Not one or the other, but both working in unison. And to get there, we need to compress time, build trust, and reimagine the stack as a system of services: intelligently connected and business-aligned.

The Cost of Tool Sprawl


The Hidden Time Drain

Users Spend 1/3rd of Time Jumping Between Tools

Despite the time overhead of employing multiple levers, we're locked in consequence of fatal data discrepancies.

01. Tool switching leads to significant productivity loss.
02. Teams struggle with inconsistent data sources and answers.
03. **Roadmap:** Unified interface. Access all capabilities across the tool stack through a single interface.

Users waste significant time switching tools. About **40%** of respondents reported spending over **30%** of their time switching tools. This inefficiency leads to confusion, resulting in three different answers to the same data question.



“


A lot of data teams run into the problem of creating metrics within their BI tool, or within something like HubSpot, or maybe another marketing tool like Google Analytics. And then, you have three different answers for the same question.

It just makes stakeholders lose trust in the data because they don't understand why there are three different answers to the same question. They don't know what to depend on. So, by shifting all of this logic left, you can really create one central source of truth.

Madison Schott, Sr. Analytics Engineer
on *Analytics Heroes by Modern Data 101*


What is the costliest aspect of maintaining a data infrastructure?

About **40%** of respondents say maintaining integrations between various data tools leads to the highest costs. Interestingly, **this is also where the respondents spend a lot of their time.**



Impact: Cost and Time

Integration inefficiencies lead to **higher costs and wasted time**




Navigating Tooling Overwhelm

The Need for Simplicity

Over **70%** of respondents indicate they use **5-7 tools or work with multiple vendors**, while **10%** use **more than 10 tools**; reflecting the increasing complexity of the Data Landscape. This reliance can lead to inefficiencies and increased costs for teams managing diverse tech stacks. Moreover, as the number of tools grows, customer experiences may decline due to fragmented processes.

To address these challenges, companies must focus on building a Data Product Hub as a single source of truth. This solution simplifies operations while enhancing visibility and governance across the entire stack. While Data Developer Platforms simplify operations through better integration and visibility across the tech stack.



Challenges

Tool Overload

The reliance on multiple tools leads to inefficiencies and increases costs, making it difficult for teams to effectively manage their operations.

Complexity Costs

Increased complexity can result in a decline in customer experiences, as teams struggle to coordinate across various tools & platforms efficiently.

“

The Unified Data Product: Breaking the Cycle

To truly break the endless cycle of rework and firefighting that plagues our data initiatives, I believe we need a fundamental shift in how we perceive and manage data. It's time we start treating data like products, not just code. For me, this means adopting a product mindset, complete with clear ownership, defined service level agreements, and user-centric design. Just as with any well-designed product, data should be intuitive and easily discoverable for its users.

The challenge I see today is that finding the right data often feels like an impossible quest. We're sifting through fragmented systems and incomplete documentation, trying to piece together a coherent picture. This is precisely why we need unified data catalogs that go beyond mere schemas. These catalogs, in my view, should provide a complete vertical slice of information, bringing together pipelines, models, governance policies, and crucial business context: all in one accessible place.

By bringing all these elements together, **we create a unified view of our data assets. This single source of truth empowers users to quickly understand** the lineage, quality, and business relevance of data, eliminating the guesswork and inefficiency that currently hinder progress. When data is treated as a product and presented in a cohesive, easily navigable manner, I'm convinced we can finally move beyond reactive firefighting and towards proactive, impactful data utilization.

Vikram Sharma

Head of Engineering, CSW at
Global Innovation Center,
Lenovo



Data Complexity



Over 70% rely on multiple tools

Most teams depend on **5-7 different** tools for various tasks, increasing complexity.



10% use over ten tools

A significant portion of users navigate through more than **10 different tools** each day.



Cost implications rise

Increased reliance on multiple tools leads to **higher operational costs** for businesses.



Efficiency declines noted

Teams struggle with productivity due to **tooling overload** and inefficient processes.



Need for a unified approach

Building a **Data Product Hub** can streamline operations and improve user experience.

73%

Chose to Integrate Disparate Tools

73% of respondents have had their data infra set up by integration of different existing tools in the market (e.g.: MDS or open source tech).




70%

Consequently Demand Automation to Manage Complexity

70% now require automation to effectively handle integration complexities.

How was your last data infrastructure set up




I feel the need for automation to manage the integrations between countless data tools

68.25%

~ 70%

Need

of the respondents feel automation is the way forward for an integrated tool stack.



“

The chasm between desired business outcome and realizable value is often widened by poor quality data sitting **across siloed systems**. Lack of real-time access to high quality data curtails innovation and often leads to poor decision making.

”



Amit Kulkarni
CTO and Co-founder,
Mojro

How would you rank the challenges in achieving a unified data experience?

85% say integrating tools across the data stack is among their **top three challenges**,

01. 85% Integration across tool stacks
02. 77% Challenge in metadata consolidation
03. 50% Lack of observability

Interestingly, challenges around both metadata consolidation and proactive observability are a **consequence of tool sprawl** where data, resources, logic, and metadata are spread across the stack with poor cross-communication bridges (poorly maintained integrations or lack of continuous updates).

Integration across tools/stack


Challenge in metadata consolidation

Managing different run-times for different tools

Lack of observability

Increased operational and maintenance cost

Others



*data (numbers) in percentage

The Productivity Drain

Time Allocation for Key Tasks



70% of respondents spend the most time on **understanding business requirements**. Spending most of your time understanding business requirements signals a fundamental breakdown in communication, alignment, or process. Here's why that's problematic:


1. Poor Upstream Clarity

If 70% of respondents are spending most of their time just trying to understand requirements, it likely means requirements are vague, inconsistent, or incomplete when handed off. This points to poor collaboration between business stakeholders and data/tech teams.

2. Waste of Expertise


Data and engineering professionals are trained to model, build, optimize, and ship solutions; not to untangle unclear intentions. When their primary job becomes “deciphering,” it underutilizes their core strengths and delays value delivery.

> continued to #3



Where do you spend most of your time?

Time Allocation for Key Tasks



Comparing time spent on critical tasks

> continued from #2

3. Delays and Rework

Misunderstood or evolving requirements often lead to inefficient development cycles, rework, or even failed projects. This directly impacts speed to insight, team morale, and trust between departments.

4. Signals Lack of a Product Mindset

In mature organizations with a data product or agile mindset, requirements are co-developed through collaborative, iterative discovery instead of being handed off as static documents. This stat shows many teams are still operating in a project mindset with weak feedback loops.

5. Hidden Cost Center

Every hour spent trying to decode business intent is an hour not spent delivering business value. Multiply that across teams and months, and you get a massive hidden cost in lost productivity.

Top 3 Time Sinks

Cumulatively, these aspects are **among the top three most time-consuming tasks**:


01. 92%
Understanding Business Requirements

02. 67%
Identifying & connecting to data sources

03. 46%
Transforming data & building ELT/ETL

The second aspect highlights a deep-rooted infrastructure gap. Teams often lack visibility into where data lives and how to access it efficiently. Time spent on data transformation and ELT/ETL signals the ongoing struggle with messy, siloed data. Instead of building insights, teams are stuck cleaning and stitching data together.

Rank the tasks where you spend your most time to least




*data (numbers) in percentage

Access to Right Data

At the Right Time

Rank the following in order of how challenging they are in terms of effort/time (top 3 ranks displayed)



The **overwhelming** statistic shows that **62% of professionals** struggle to obtain the right data when needed. This highlights the critical need for efficient data management systems to streamline access and enhance productivity.

“

Finding the right data shouldn't feel like guessing a password. We need catalogs that don't stop at schemas. They should show the full stack. Pipelines, Models, Governance & Business Context. So people can actually use the data.

”




Surendra Mathe
Partner, **Gensler**

Immediate Data Access Remains Out of Reach

Business experts are looking to get data access immediately, but for more than **66%** of respondents, access takes more than a day and for **20%**, more than five days (often when the need is over).

The traditional path from data to insight is paved with good intentions and long delays. Data is pulled, cleaned, transformed, modelled, validated, packaged, and then (maybe) shared. Each step is owned by a different person, team, or tool. Business teams wait. Questions pile up. By the time answers arrive, the opportunity has passed or the question has changed. The distance between curiosity and clarity has become the silent killer of momentum.


How often do you need to rework your data products (or tables/dashboards) because they are not meeting business needs



Shrinking that distance.

Modular data products have proven to shrink the path from data to insights, since it starts directly from user requirements and operates wholly within that boundary. With clearly defined owners, analysts are embedded directly into business units, and semantic layer brings standardization to ensure consistent definitions across tools.

They automate lineage tracking and alerting to catch upstream issues early, and enable self-serve interfaces so business users could explore trusted datasets without waiting on data teams. Each step narrows the distance from request to insight, and consequently from insight to action. Turning what used to take days into minutes.



PRE DATA USAGE TIME-SPEND ANALYSIS


01. **63%** spend 1/5th or more of their time on **maintenance**.


02. **65%** spend 1/5th or more of their time **figuring out what data should be used** to implement any new data project or product.

If teams are spending over 1/5th of their time just finding data or maintaining it, it is evident how the path to data usage is a distant cry. Valuable hours are lost fixing and digging through dashboards, Slack threads, old Confluence pages, or tribal knowledge, none of which scale.

A unified interface like a data product hub changes this. It offers a single place to search, explore, and understand available data assets, complete with context, ownership, quality signals, and documentation (maintenance delegated to underlying data developer platform). Instead of chasing scattered sources, teams can focus on using data, not finding it.

Understanding the impact of maintenance and data discovery on productivity






Top Post-Deployment Challenges

UNDERSTANDING THE THREE MAIN HURDLES DATA PROJECTS FACE

The chart illustrates that almost **60%** of **respondents** identified the inability of data projects to keep pace with business. These respondents feel data projects are too slow with limited business value and do not meet or respond slowly to business needs.

Data Teams Misaligned: Rework Data & Data Products Regularly

How often do you need to rework your data products (or tables/dashboards) because they are not meeting business needs



60% of Data Teams rework their tracks often

Over 13% rework it more than 25% of the time. Not due to incompetence, but because of a deeper systemic issue. The root cause lies in a reactive culture, where business questions evolve faster than data teams can anticipate, leading to misaligned requirements, undocumented logic, and fragile pipelines. Data teams are forced into a constant state of firefighting.

Breaking the Cycle

Breaking the cycle begins with treating data products like real products. That implies investing in clear ownership, reusable models, standardized definitions, and a product-oriented mindset where collaboration with business stakeholders happens upfront instead of after delivery.

A “model-first” approach, where semantic clarity and use-case alignment are prioritized from the beginning, can prevent cascading errors and reduce ambiguity.

Data-to-Insights

Long Meandering Path



Time Spent Understanding Requirements

~**70%** spend the most time on understanding business requirements.

Time Spent Identifying & Connecting to Data

62% struggle to obtain the right data when needed.

Waiting on Access Post-Identification

For **66%**, access takes over a day and for 20%, more than 5 days (often when the need is over).

Insights After Business Moves On

~**60%** cite time as top two post-deployment challenges: 1/ data projects too slow with limited business value, 2/ not meeting or slow response to needs.

Closing the Gap

While time sinks are inevitable, major sinkholes are avoidable.

The respondents have themselves come up with

the data stack they need.

CLOSING THE GAPS

The Data Stack Teams Need

SIMPLIFYING COMPLEXITY & BECOMING USER-DRIVEN

Way Forward: Automation

Nearly 70% of data professionals report feeling the need for automation to manage integrations between multiple tools: a clear signal that the current landscape of fragmented platforms is unsustainable. As data stacks grow more modular, data citizens need the freedom to focus on outcomes, not pipelines. They require automation that ensures seamless data flow across systems, without manual stitching, fragile dependencies, or constant engineering intervention.


They’re NOT asking for more tools, but less friction to manage existing ones.

They want a layer that connects their tools, handles the handoffs, and guarantees data reliability without needing to open a terminal or wait for another sync job to be debugged. This need is most acutely felt at the foundation: **85%** of architects and **75%** of infrastructure professionals highlight it, pointing to a significant **gap in the data infrastructure itself**. On the other end, **75%** of data app developers (responsible for the bridge between data and real business impact) also feel the pain.


The answer lies in a Unified Utility Plane

A self-serve abstraction layer that automates tool-to-tool communication and integrates seamlessly with the broader data ecosystem. Platforms built per the Data Developer Platform Standard enable this by acting as a singular consumption layer, abstracting away complexity of integration maintenance and orchestration. Instead of maintaining integrations, teams simply declare what data they need and where it needs to go, the self-serve platform handles the rest. With built-in intelligence and observability, integrations become invisible yet resilient.

I feel the need for automation to manage the integrations between tools



The Gap Radar



What capability do you feel is lacking in traditional or current data stacks?
*multiple choice

Native Integration Unity Over Sprawl

This radar highlights the most in-demand capabilities that data leaders believe are lacking in traditional or current data stacks. Semantic layers and centralized governance with decentralized policy execution are both cited as lacking by **over 50%**. This reflects a clear need for standardized, business-aligned data interpretation and policy enforcement frameworks that can operate flexibly across domains without compromising control.

Another critical gap, as cited by **nearly 70%**, is the absence of catalogs that go beyond just data cataloging, extending to infrastructure and services. This signals a demand for holistic visibility into not just datasets, but the entire ecosystem powering them (pipelines, APIs, models, and compute environments). Additionally, **62%** point to natively integrated data quality as a major missing piece, highlighting frustration with the fragmented, bolt-on nature of current quality checks.

Interestingly, lakehouses and “others” rank significantly lower, suggesting that while storage architecture matters, it’s no longer the core pain point. Instead, the future-proof data stack must focus on interoperability, governance, and trust with abstraction layers that simplify complexity, and unified platforms that integrate intelligence directly into the data lifecycle.

Closer Look Into Data Quality & Observability

Data quality isn't just a technical concern, it's a foundational and cultural gap (pointing to trust) in the modern data stack. **62%** of respondents identified the lack of natively integrated data quality as the biggest missing capability in today's setups. The bolt-on approach to quality checks, often stitched together through fragmented tooling, simply **doesn't scale**. This absence leaves teams vulnerable to inconsistent data, delayed insights, and a **rising cost of ownership**.

The call is clear: data quality needs to be **simpler and self-served**. An overwhelming **61%** of practitioners believe that a self-serve platform for declarative data quality (or **low-code or no-code**) is the most effective way to manage quality checks and service level objectives (SLOs). At the same time, **50%** want native observability built into the data platform itself, to reduce the **burden of setting up monitoring or relying on external alerting systems**.

Taken together, these numbers reveal a broad dissatisfaction with the status quo and a growing appetite for **automation and abstraction**.

Way forward as indicated by responses

Instead of stitching quality, observability, and governance post-hoc, data platforms must treat them as first-class citizens: **embedded** into the data product lifecycle. A data product approach ensures that every dataset or pipeline comes with quality, reliability, and monitoring built-in by default.

Combined with self-serve infrastructure for **declarative quality operations**, this model empowers users to define, monitor, and improve data quality without relying on specialized engineering support, closing the gap between operations and outcomes.

Effective observability starts with clear specification from ideation itself, a capability ensured by requirements-first data products. Whether it's **data** (e.g., freshness, volume, schema drift), **infrastructure** (e.g., pipeline latency, service uptime), **code** (e.g., test failures, logic errors), or **resources** (e.g., compute usage, storage thresholds). Without this **clarity across the “vertical slice”**, monitoring remains shallow and reactive, often relying on post-hoc alerts instead of proactive, built-in safeguards.

“

Native observability across all your enterprise data, at rest and in motion, isn't optional, it's survival. When you are making hundreds of business decisions using data, reactive cognizance of data issues means you're already too late.


”



Mukund Raghunath

CEO and Founder,
Acies


What capability do you feel is lacking in traditional or current data stacks



6 in 10 Demand Native Data Quality

62% point to lack of natively integrate quality, highlighting issues with the fragmented, bolt-on nature of quality checks.


What might be the best way to reduce the workload to define quality checks and SLOs



6 in 10 Demand Declarative or Self-Served Quality

61% of the respondents identified the need for a Low-Code/No-Code approach to manage Data Quality.

What might be the best way to set up observability (monitoring & alerting)



50% Feel the Need for Native Observability

50% of the respondents feel the need for a data platform with native observability to simplify monitoring.

Closer Look Into Data Discovery & Understandability

As modern data estates grow in scale and complexity, the traditional definition of a "data catalog" is becoming increasingly outdated. It's no longer enough to just list datasets and their schemas. Nearly **70%** of respondents express a clear need for catalogs to evolve, moving beyond raw data to also include infrastructure components, logic, services, pipelines, models, and access policies. What they're really asking for is not just a catalog, but **a map: a vertically integrated, navigable view** of the entire ecosystem surrounding their data assets.

This desire isn't academic. It's practical. Teams want to trace a broken dashboard to its upstream model, understand what pipeline powers a table, and know which team owns which transformation logic. In other words, they seek **understandability that spans beyond the "what" of data to the "how," "where," and "who."** Without this, discoverability becomes superficial. Users may find a table but remain blind to whether it's trustworthy, relevant, or reproducible. Not know how and how fast to get it fixed.

40% of respondents further point to a deeper challenge: the need for automatic end-to-end curation of information across the stack. The current manual processes (documenting schemas, dependencies, or lineage) don't scale and quickly go stale.


What's missing is a **system that understands the interconnected fabric of the data stack**, curates insights in real-time, and makes them accessible by default. This is no longer a "nice-to-have"; it's essential for a data-driven organization.

When data is packaged as a product, it comes with its own contract, ownership, documentation, and quality guarantees. But the true transformation happens when these products are built on a self-serve platform that **automates** metadata capture, policy enforcement, and service integration.

By abstracting away the complexity and surfacing context-rich, automatically curated products, organizations empower data consumers to confidently navigate, understand, and act on data **without requiring tribal knowledge or heroic effort**.

A modern self-serve infrastructure doesn't just make discovery possible; it makes it a culture. It turns every action, every model deployment, pipeline update, or schema change, into metadata that enhances the discoverability of the entire system. That's what real understandability looks like: living documentation instead of static metadata.


What capability do you feel is lacking in traditional or current data stacks



~ 70% Need Catalogs to Go Beyond Data

7 in 10 feel the lack of a catalog that goes beyond data and also catalogs data infrastructure & services. Implying a need for domains to have a “vertical view” of the elements associated with data assets.

What might be the best way to explore existing data assets before starting actual work on them




4 in 10 Demand Automatic E2E Curation

40% feel the need for a new Data Catalog that automatically curates information from the entire technology stack for data exploration (discoverability, understandability, & accessibility).

Closer Look Into Optimal Governance Framework

As data volumes grow and responsibilities fragment across domains, the debate between centralized and decentralized governance intensifies. **Centralized governance** promises consistency, security, and compliance but often at the cost of agility and domain-level autonomy. **Decentralized governance** offers speed, flexibility, and local ownership but it can lead to chaos, silos, and fragmented policies. The reality is, **most organizations need both**. In fact, **almost 60% of respondents advocate for a hybrid model**: centralized governance for control and alignment, with decentralized policy execution at the domain level for agility and relevance. This balance allows organizations to maintain a strong backbone of standards while empowering teams to **make contextual decisions**.

Implementing this hybrid governance model effectively requires more than just organizational intent: it requires a strong technical foundation. That's where self-serve data platforms, such as data platforms built in alignment with the Data Developer Platform (DDP) Standard, come into play. These platforms are designed to embed governance as a native capability, not an afterthought. They provide **centralized policy definition** (think: access controls, data quality SLOs, auditability) while enabling **decentralized enforcement** through composable building blocks, automation, and domain-specific tooling. With such a system, governance scales not by gatekeeping, but by empowering domain teams with guardrails: turning compliance from a bottleneck into a seamless, integrated part of the data lifecycle.



Cumulative View

Signalling a Simpler Interface

50%

feel the lack of a semantic layer in traditional or current data stacks standardized for business-aligned data interpretation that enable data accessibility flexibly across domains

15%

demand to model existing, known data assets into a logical query interface for data exploration.

20%

identified the need for a semantic layer for data exploration.

13%

need an intuitive query interface that converts Natural Language sentences into appropriate SQL statements for data exploration.

“

Business users don't think in tables and joins, they think in metrics and dimensions. A semantic layer today is a necessity to bridge the gap between technical implementation and business understanding.

”



Mukund Raghunath

CEO and Founder,
Acies


The Twin Engine Driving Acceleration for Data Teams

Cumulatively, **65%** cite having data products and the right data models as the top effective and essential capabilities to accelerate their work associated with data consumption.

In an isolated view, having the right data models takes the first place as the most effective solution with **71%** ranking it among the top two effective options. This is closely followed by Data Products with **60%** citing it among the top two solutions for accelerating data consumption.



As a data consumer, arrange the following capabilities based on how effective they can be in accelerating your work



*data (numbers) in percentage

What Our Data Platform Builders Are Saying

The survey results shine a light on the upcoming heavy tax of being ‘almost AI-ready.’ The data platform approach should be collapsing that gap. Must ensure production-grade assets **rich with context** that trigger and accelerate AI adoption. It is the new normal. AI devours data. This is THE moment when we need a better data foundation than ever before.

This validated a painful reality: today’s data workflows, especially in enterprises, aren’t built for the AI tipping point. AI needs a data architecture that’s heavily context-centric. It brings structure and adds a semantic spine to data ecosystems, giving teams a productised and AI-ready foundation. Not pipelines, but think utility products **built for** AI Agents (machine users).

The clearest signal from this survey? The market is done with duct-taped solutions. Enterprises want to move from data experimentation to AI execution. Fast. Platforms like DataOS are purpose-built to accelerate that leap, enabling teams to shift from managing pipelines to delivering trusted, AI-ready outcomes.



Animesh Kumar
CTO and Co-Founder,
The Modern Data Company



Travis Thompson
Chief Architect, DataOS
The Modern Data Company



Sagar Paul
SVP, Global Growth & Solutions,
The Modern Data Company



About The Modern Data Company

The Modern Data Company was founded to radically simplify how organizations manage, access, and interact with data. Modern's data operating system, DataOS®, unifies all data across silos, at scale. The composable layers of DataOS evolve with your existing technology stack to future-proof your data initiatives.

themoderndatacompany.com info@tmdc.io



About Modern Data 101

Modern Data 101 is a publication and community for people building data platforms, designing data teams, and architecting the invisible. In a world of countless tools, trends, and templated thought leadership, Modern Data 101 slows things down. Every community-driven publication is a lens to data and AI architectures, semantics, and organisational design.

moderndata101.com

community@moderndata101.com

