

Shifting from Point Solutions to Unified AI-Native Data Platforms

Over 70% of data teams work across 5 to 7 tools, spending 30%+ of their time moving between disconnected interfaces

Table Of Contents

1. Executive Summary
2. The Industry Landscape: Convergence From Every Direction
3. From Niche to Platform: The Story of Some of the Popular Platforms
4. Weighing Tooling Siloes and Platform Consolidations
5. Why the Tool Economy Is Losing Ground
6. How Platforms Are Becoming AI Native
7. What a Unified Platform Is Made Of
8. Where Convergence Stops
9. Platform Evaluation: 5 Criteria to Distinguish True Platforms
10. Conclusion
11. References

Executive Summary

From The Modern Data 101 Community

This report maps how Snowflake, Databricks, dbt and DataOS, and the wider industry around them, are converging on one governed substrate. Choose this if you want the quote to name the examples and signal the industry view.

Five years ago Snowflake sold a cloud data warehouse, Databricks sold the lakehouse, and dbt sold a SQL transformation tool. Each stopped where its category stopped, and buyers assembled the rest. In 2026 all three sell platforms that span ingestion, governance, semantics and AI, and DataOS, built as an operating layer over existing stacks, is extending toward the same destination. The expansion follows one logic: the tool economy depended on integration that buyers had to do themselves, and platforms now do that integration for them, with governance, context and agents included.

The scale of the shift shows in the numbers.

Databricks moved from \$600 million ARR in 2021 to a \$7 billion revenue run-rate in 2026, and dbt, now merged with Fivetran, serves 100,000 data teams.

What this paper argues

- Four platforms expanded from niche categories into whole-stack platforms, because buyers reward the removal of joins.
- The tool economy loses ground for plain reasons: the integration tax compounds, AI consumes shared context, platforms match specialists on commodity layers, and open standards bring openness inside the platform.
- AI nativity is the next layer of expansion. Context, agent interfaces, agent governance and operational data are being built into the platform itself.

The Industry Landscape: Convergence From Every Direction

Convergence in data is an industry-wide realignment, and it is not one company's strategy. The modern data stack grew on a simple bargain: pick the best tool for each layer and accept the work of joining them. The landscape swelled to 2,011 logos in 2024 on the strength of that bargain. Buyers now report the cost of keeping it. Over 70% of practitioners work with more than five to seven tools or three to five vendors, and about 40% spend over 30% of their time moving between them.

Vendors of very different origins have responded in the same direction. Cloud hyperscalers, data cloud platforms, enterprise application suites, pipeline vendors and observability specialists are all moving toward a single governed substrate, and capital is following. DevOps and DataOps deal value reached \$99.8 billion in the twelve months to mid-2026, 2.1 times the prior year. Demand is just as clear: 77% of practitioners call a converged platform valuable.

The rest of this paper uses four platforms as examples of this movement. They are illustrations of a market-wide pattern, not the whole of it.

Roughly 5% of companies qualify as genuinely AI-driven, another 35% are still catching up

From Niche to Platform: The Story of Some of the Popular Platforms

Around 2021 each of these four started from a well-defined job. By 2026 each owns most of the stack around that job.

Snowflake's growth tracks its scope. Product revenue grew eight times, while customers spending over \$1 million grew from 77 to 733, about 9.5 times, against a 3.2-times rise in total customers, from 4,139 to 13,328 customers. Customers that adopt more of the platform spend disproportionately more, which is the commercial engine behind the expansion. Along the way Snowflake added ingestion, operational Postgres, an open catalogue, AI, agents and observability.

Databricks created the lakehouse category and was valued at \$38 billion in 2021. By August 2026 its run-rate had passed \$7 billion at 80% growth. Warehousing alone is now a \$1.5 billion line, AI products passed a \$1.4 billion run-rate in February 2026, and Lakebase, its operational database, passed \$100 million. A company that sold one architecture now sells the whole stack around it.

By August 2026 its run-rate had passed \$7 billion at 80% growth.

dbt began as the transformation step in the middle of the stack. More than 80,000 data teams used it by October 2025, and the completed Fivetran merger now spans movement, transformation and semantics. 3,000 organisations already run dbt on Databricks, which shows how tightly the platforms and the former tools now overlap.

The expansion path is the same in all four

Each platform started with one job, absorbed the adjacent layers that created the most joins, made governance and the catalogue central, added a semantic layer, and then opened an agent interface. The order is consistent because each step depends on the one before: governance needs the data in one place, semantics need governance, and agents need semantics.

3,000 organisations already run dbt on Databricks, which shows how tightly the platforms and the former tools now overlap.

Platform	Starting point	Scope in 2026
Snowflake	Cloud data warehouse: storage, compute and SQL	AI Data Cloud: managed ingestion (Openflow), operational Postgres, an open catalogue and governance layer (Horizon), semantic context, AI and agents (Cortex, CoWork, CoCo) and observability
Databricks	The lakehouse: Spark, Delta Lake and machine learning on open storage	Data and AI platform: warehousing, Lakeflow pipelines, Unity Catalog governance, Lakebase operational Postgres, Genie and Agent Bricks, a security lakehouse and a customer data platform
dbt	Open-source SQL transformation and the analytics engineering workflow	With Fivetran: ingestion, transformation, activation and semantic modelling, the Fusion engine with a metadata layer, and agent tooling through a remote MCP server
DataOS	A data operating system layered over existing stacks, with the data product as its unit	Policy-as-code governance, metadata and lineage, a unified semantic layer, observability, and MCP-native data products for agents

Image: Four platforms, then and now: where each started and what each covers in 2026 | Source: Authors

Weighing Tooling Siloes and Platform Consolidations

Platforms win because they remove the cost the tool economy was built on. The modern data stack asked buyers to pay an integration tax in exchange for a specialisation premium: the best tool for every layer, replaceable at will. The landscape swelled to 2,011 logos in 2024 on the strength of that bargain, and the bill arrived with it. Over 70% of practitioners work with more than five to seven tools or three to five vendors, about 40% of practitioners spend over 30% of their time moving between tools, and upkeep adds 15 to 18 percent on top of licence spend.

The tax compounds, because each new tool adds a join to every tool already deployed, with its own connector, its own access model and its own definition of a customer. A platform that internalises those joins turns the tax into a product feature, and buyers pay for the feature gladly. That is the logic behind the revenue pattern discussed in previous sections.

AI raised the stakes. Humans were the integration layer of the modern data stack: they reconciled the conflicting number and delayed a decision until someone confirmed the data. An agent cannot do that. The average enterprise now runs 12 AI agents, half of them in silos, and almost 70% of practitioners say their data is not yet clean or trustworthy enough for AI. When the consumer was a person, a join cost minutes of judgement. When the consumer is an agent, a join costs correctness.

Lower cost and one answer

42% of data leaders already rank simplifying the data architecture among their top priorities, and 80% of practitioners rank a semantic layer with standardised definitions as the top AI enabler. Shared definitions end the conflicting numbers that erode trust.

Governance that applies itself

81% of practitioners put consistent governance among their top requirements, 84% want unified access, and 72% want self-service without waiting on IT. One policy engine turns governance into a property of the platform.

Demand is already there

Practitioners ask for convergence rather than more tools: 77% of practitioners call a converged platform valuable, and nearly 40% rate it extremely valuable. The platforms that moved first are the ones answering that demand.

Why the Tool Economy Is Losing Ground

The multi-tool ecosystem offered four advantages: depth of specialisation, pace of innovation, flexibility to swap parts, and fit to purpose. Each is fading for a reason that needs little explanation.

Advantage	Why it is fading	What replaces it
Depth of specialisation	Platforms now ship good-enough storage, query, quality and lineage, and storage, scalability and query performance have become table stakes	Depth where it still matters, delivered inside the platform or plugged into it
Pace of innovation	Platforms keep pace by building and buying, and consolidation compresses a specialist category before it matures	Faster releases of integrated capability
Flexibility to swap parts	Swapping a part means re-wiring every join it touches, while open standards now deliver flexibility inside a platform	Swappable components behind stable interfaces
Fit to purpose	Each purpose-built tool re-implements governance, catalogue and policy for itself	Purpose-specific engines on one governance substrate

Image: The four advantages of the tool economy, why each is fading, and what replaces it | Source: Authors

Two further reasons sit underneath. Governance cannot be bolted on tool by tool: the leading obstacles to aligning data and AI strategy are separate governance models (41%), and the lack of unified platforms (37%). And standards have cheapened the join. The classic argument for a fragmented stack was that only independent tools could stay open.

Open table formats and agent protocols changed that, with 10,000 active public servers now speaking the Model Context Protocol and semantic standards following. A platform can embrace outside tools through standard interfaces, so openness no longer requires fragmentation.

The market is acting on these reasons. DevOps and DataOps deal value reached \$99.8 billion in the twelve months to mid-2026, 2.1 times the prior year.

DevOps and DataOps deal value reached \$99.8 billion in the twelve months to mid-2026, 2.1 times the prior year.

Move	Layer	Status	What it changes for the buyer
Salesforce + Informatica	Integration, catalogue, governance, master data	Closed Nov 2025	A data control plane enters an application suite
IBM + Confluent	Streaming and data in motion	Closed Mar 2026	Real-time data joins hybrid cloud and AI portfolios under one owner
Fivetran + dbt Labs (with Census, Tobiko)	Ingestion, transformation, activation	Closed Jun 2026	One vendor spans the pipeline from source to semantic model
Databricks + Tabular, Neon; Snowflake + Crunchy Data	Table formats, transactional Postgres	2024 to 2025	Open formats and operational databases move inside the platform
Databricks LTAP, Lakehouse//RT, CustomerLake	Transactional storage, real-time serving, customer data	Announced Jun 2026	Real-time serving runs natively in one governance model, with no separate permissions layer or sync pipelines
Datadog + Metaplane; Soda + nannyML; Snowflake + Observe	Data and ML observability	2025 to 2026	Observability becomes a native capability of larger platforms
SAP + Dremio, Prior Labs	Lakehouse query, AI	Announced May 2026	An application vendor builds toward an AI data control plane

Image: The consolidation record: acquisitions and launches pulling data layers into larger platforms | Source: Authors

The consolidation we flagged in October 2025, when the Fivetran and dbt merger was still a rumour, has arrived. The work now is making it architectural.

6

How Platforms Are Becoming AI Native

AI native means agents are first-class consumers of the platform. Four capabilities mark it: a context layer, an agent interface, governance built for agents, and operational data that agents can act on. All four platforms are building all four.

Capability	Snowflake	Databricks	dbt	DataOS
Context layer	Horizon Context and Cortex Sense, a governed semantic layer inside the catalogue	Genie Ontology and Unity Catalog metadata and lineage	Fusion metadata layer generated at compile time; MetricFlow open sourced for semantic interchange	A unified semantic layer built into every data product
Agent interface	CoWork, CoCo, Managed Agents, and an announced acquisition of an MCP platform	Genie and Agent Bricks	Remote MCP server and dbt agents	Data products exposed through MCP, REST, GraphQL and SQL
Governance for agents	Agent Identity and runtime AI guardrails	Unity AI Gateway for multi-AI governance and cost control	Column-level lineage and impact analysis before anything runs	Policy-as-code governance and zero-trust access
Operational data	Snowflake Postgres and Openflow ingestion	Lakebase and LTAP, built for agents	Fivetran data movement feeding governed models	Queries and activates data across sources without copying it

Image: What “AI native” looks like: four capabilities across the four platforms | Source: Authors

AI is already a revenue line. Databricks' AI products passed a \$1.4 billion run-rate in February 2026, and Snowflake posted 37% product revenue growth in its second fiscal quarter of 2027, driven by a step-up in AI revenue alongside the core platform.

The contest is over context. An agent cannot reconcile two definitions of revenue, so the platform has to hold one. Semantic layers are on course for critical-infrastructure status by 2030, and 60% of agentic analytics projects that rely on the Model Context Protocol alone are expected to fail by 2028 without a consistent semantic layer underneath. Half of agent deployment failures by 2030 are expected to trace to insufficient runtime governance and multisystem interoperability. Whoever holds the definitions holds the agent's view of truth, which is why every platform in this paper is building its own context layer and why semantic standards matter.

What AI native adds to convergence

Convergence began as a cost argument: fewer joins mean less integration work. AI nativity makes it a correctness argument. A platform that holds data, policy, definitions and agent access in one substrate can guarantee that an agent sees what a person sees, and can enforce the same rules on both. That guarantee cannot be assembled from parts after the fact.

What a Unified Platform Is Made Of

A unified platform is easy to describe and hard to fake. It has one substrate that every layer reads from and writes to, so that a definition, a permission or a lineage record is created once and applies everywhere. The table lists the layers, what becomes shared in each, what the buyer gets, and how far each has converged.

Layer	What becomes shared	What the buyer gets	Convergence status
Storage and compute	One copy of data in open table formats, served to analytics, operations and AI	No sync pipelines between systems	Converging on open formats
Ingestion and transformation	Tested transformations and one lineage graph from source to model	Fewer connectors and one place to debug	Largely converged
Streaming	Real-time and batch data on one governed substrate	One pipeline for fresh and historical data	Converging
Metadata and lineage	One catalogue and one lineage graph across every layer	Impact analysis and root cause in one place	Converging fast
Governance and policy	One policy engine enforced in pipelines, at query time and on agent access	A rule written once applies everywhere	Converging fast
Semantic layer	One set of versioned business definitions every tool and agent can read	One answer to the same question	Contested and accelerating
Orchestration and observability	One scheduler and one health signal across pipelines, quality and models	Faster detection and automated remediation	Largely converged
Interface layer	Standard protocols for people, applications and agents	Specialists plug in without reintegration	Standardising

Image: The layers of a unified platform, what each shares, and how far each has converged | Source: Authors

60% of agentic analytics projects that rely on the Model Context Protocol alone are expected to fail by 2028

The sequence is consistent across platforms. Storage converges first, because open table formats make one copy practical. Metadata and policy converge next, because they need context from every layer. The semantic layer and the agent interface converge last, because they depend on everything beneath them. As layers converge, the scarce asset becomes the context beneath them: the lineage graph, the semantic model and the policy engine that every layer reads from. Salesforce's Informatica purchase and SAP's planned Dremio and Prior Labs acquisitions chase the same prize, a data control plane, and half of organisations are expected to use agents to turn policy into machine-verifiable data contracts by 2030.

8

Where Convergence Stops

An industry view also has to say where the pattern does not hold. Convergence is strong, but it is not total.

- Latency-critical and sovereign workloads keep purpose-built components, and sector-specific designs continue to dominate regulated workloads.
- Regional rules fragment the map. 35% of countries are expected to sit on region-specific AI platforms by 2027, so some multinationals will run more than one platform.
- Specialists survive where they own distinctive context. A tool that holds a unique lineage graph, semantic model or policy asset and speaks open interfaces has a durable place. A tool that cannot say what context it owns becomes a feature inside someone else's platform.
- Lock-in is the main cost of convergence. Only 8% of executives list it as a challenge today, because its cost arrives at exit.
- The realistic end state is not one vendor for everything. It is a converged substrate with a short list of specialists plugged in through open interfaces.

Platform Evaluation: 5 Criteria to Distinguish True Platforms

Convergence pays when it is architectural and the interfaces stay open. Five checks separate a unified platform from a bundle of products behind one invoice.

Check	The question to ask	What good looks like	What to avoid
One platform	Does a policy written once apply to every layer, or is it recreated per module?	A single policy engine, with enforcement logs showing the same policy across ingestion, query and agent access	A separate permission model for each acquired product
One source of truth	Where do metric definitions live, and can an outside tool read them?	Open, versioned definitions exposed through a standard interface	Definitions that only work inside one vendor's BI
Open	Can data and metadata leave in open formats at contract end, at no extra cost?	Open table formats, with exportable lineage, policies and definitions	An open format managed behind a proprietary catalogue
AI-ready	Can an agent discover, query and be audited through a standard protocol?	A standard protocol with identity and logs on every call	Agent features reachable only through a vendor SDK
Integrated	Was it built or bought, and has the join actually gone?	One codebase and one metadata store	Two stores behind one invoice

Image: Five checks that separate a unified platform from a bundle of products behind one invoice | Source: Authors

Convergence has edges worth planning for. Latency-critical streaming and sovereign deployments keep some purpose-built components, sector-specific designs continue to dominate regulated workloads, and regional rules will push some multinationals to run more than one platform, with 35% of countries expected to sit on region-specific AI platforms by 2027. The answer is a unified substrate with a short list of specialists on open interfaces.

Treat portability as a design requirement

The main risk to manage is lock-in, and it is easy to underweight: only 8% of executives list it as a challenge today, against 38% for fresh data, because its cost arrives at exit. Vendors of unified platforms list it themselves among adoption challenges. Written into the design as open formats, exportable policies and standard protocols, it becomes a contract term instead of a surprise. Start from the cost your team can already measure: the share of engineering time spent maintaining joins, which for about 40% of practitioners exceeds 30% of their working time.

Conclusion


Within two years, the argument over best-of-breed versus platform will read like the argument over on-premises versus cloud: settled in the middle and loudest at the edges. The four platforms in this paper took different roads to the same place. Snowflake came from the warehouse, Databricks from the lakehouse, dbt from the transformation layer, and DataOS from the operating layer, and all four now offer governed data, shared meaning and an agent interface in one substrate. Most enterprises will run a converged substrate with a short list of specialists that each own a distinctive context asset and speak open interfaces. Tools that cannot say which context they own will survive as features inside someone else's platform. Buyers who treat portability as a design requirement will converge on their own terms. Buyers who do not will converge on a vendor's.

References

- [1] Snowflake Reports Financial Results for Q4 and Full Year Fiscal 2021. Snowflake, Feb 2021.
<https://www.snowflake.com/en/news/press-releases/snowflake-reports-financial-results-for-the-fourth-quarter-and-full-year-of-fiscal-2021/>
- [2] Snowflake Form 10-K, fiscal year ended 31 January 2021. Snowflake, SEC filing.
<https://www.sec.gov/Archives/edgar/data/1640147/000164014721000073/snow-20210131.htm>
- [3] Snowflake Form 10-K, fiscal year ended 31 January 2026. Snowflake, SEC filing.
<https://www.sec.gov/Archives/edgar/data/1640147/000164014726000008/snow-20260131.htm>
- [4] Snowflake Reports Financial Results for Q2 Fiscal 2027. Snowflake, Sep 2026.
<https://www.sec.gov/Archives/edgar/data/0001640147/000164014726000033/fy2027q2earnings.htm>
- [5] AI agents, open data and governance take center stage at Snowflake Summit. SiliconANGLE, Jun 2026.
<https://siliconangle.com/2026/06/02/ai-agents-open-data-governance-take-center-stage-snowflake-summit/>
- [6] Snowflake Summit 2026: Product Announcement Recap. SELECT, 2026. <https://select.dev/posts/snowflake-summit-2026-product-announcement-recap>
- [7] Databricks raises \$1.6B at \$38B valuation as it blasts past \$600M ARR. TechCrunch, Aug 2021.
<https://techcrunch.com/2021/08/31/databricks-raises-1-6b-at-38b-valuation-as-it-blasts-past-600m-arr/>
- [8] Databricks surpasses \$5.4 billion revenue run-rate. Databricks, Feb 2026.
<https://databricks.com/company/newsroom/press-releases/databricks-grows-65-yoy-surpasses-5-4-billion-revenue-run-rate>
- [9] Databricks lands \$5bn funding, valuation climbs to \$190bn. Yahoo Finance, Aug 2026.
<https://finance.yahoo.com/technology/ai/articles/databricks-lands-5bn-funding-valuation-073323521.html>
- [10] Databricks Data + AI Summit 2026 recap. Flexera, Jun 2026. <https://www.flexera.com/blog/perspectives/databricks-data-ai-summit-2026/>
- [11] Databricks expands Lakehouse to unify OLAP and OLTP. Blocks and Files, Jun 2026.
<https://www.blocksandfiles.com/data-management/2026/06/18/databricks-expands-lakehouse-to-unify-olap-and-oltp/5258373>
- [12] How to Repoint dbt ETL Pipelines to Databricks. Databricks, 2026. <https://www.databricks.com/blog/how-repoint-dbt-etl-pipelines-databricks>

- [13] dbt Labs announcements at Coalesce 2025. dbt Labs, Oct 2025. <https://www.getdbt.com/blog/dbt-labs-cost-optimization-agentic-ai-product-announcements>
- [14] Fivetran + dbt Merger and Fusion Engine Explained. Kestra, Jun 2026. <https://kestra.io/resources/data/fivetran-dbt-merger-fusion-engine>
- [15] Fivetran and dbt Labs Complete Merger. BigDATAwire, Jun 2026. <https://www.hpcwire.com/bigdatawire/this-just-in/fivetran-dbt-labs-complete-merger-to-create-the-data-infrastructure-for-trusted-ai-agents/>
- [16] DataOS: Business-Ready Data Products. The Modern Data Company. <https://www.themoderndatacompany.com/dataos>
- [17] Behind Modern's operating system for data. Blocks and Files, Jul 2026. <https://www.blocksandfiles.com/data-management/2026/07/13/behind-moderns-operating-system-for-data/5270604>
- [18] Best AI-Ready Data Platforms: Data Readiness Guide. Modern Data 101, 2026. <https://www.moderndata101.com/blogs/best-ai-ready-data-platforms-data-readiness-guide>
- [19] Data Products: The Essential Context for Enterprise AI. Modern Data 101, 2026. <https://moderndata101.substack.com/p/data-products-the-essential-context>
- [20] The Modern Data Report 2026. Modern Data 101, 540+ practitioners, 2026. <https://www.moderndata101.com/modern-data-report-2026>
- [21] The Current Data Stack Is Too Complex. Modern Data 101, 230+ practitioners, Mar 2025. https://medium.com/@community_md101/the-current-data-stack-is-too-complex-70-data-leaders-practitioners-agree-b460821b07dd
- [22] The Modern Data Stack's Final Act. Modern Data 101, Oct 2025. <https://moderndata101.substack.com/p/the-modern-data-stacks-final-act>
- [23] The 2024 MAD Landscape. FirstMark, Matt Turck, 2024. <https://www.mattturck.com/mad2024>
- [24] 2026 Connectivity Benchmark Report. Salesforce (an integration vendor), Feb 2026. <https://www.salesforce.com/news/stories/connectivity-report-announcement-2026/>
- [25] IBM Acquiring Confluent to Boost AI Development Capabilities. TechTarget, Dec 2025. <https://www.techtarget.com/searchdatamanagement/news/366636098/IBM-acquiring-Confluent-to-boost-AI-development-capabilities>
- [26] Global DevOps and DataOps Report Q2 2026. Drake Star, Aug 2026. <https://www.drakestar.com/news/global-devops-and-dataops-report-q2-2026>

- [27] Gartner Top Predictions for Data and Analytics in 2026. Gartner, Mar 2026. <https://www.gartner.com/en/newsroom/press-releases/2026-03-11-gartner-announces-top-predictions-for-data-and-analytics-in-2026>
- [28] Gartner Top Predictions for IT Organizations in 2026 and Beyond. Gartner, Oct 2025. <https://www.gartner.com/en/newsroom/press-releases/2025-10-21-gartner-unveils-top-predictions-for-it-organizations-and-users-in-2026-and-beyond>
- [29] Converging Data Management Markets Force D&A Leaders to Assess Options. Gartner, Feb 2026. <https://www.gartner.com/en/documents/7460826>
- [30] Takeaways From The Forrester Wave: Data Lakehouses, Q3 2026. Forrester, 2026. <https://www.forrester.com/blogs/takeaways-from-the-forrester-wave-data-lakehouses-q3-2026/>
- [31] The Forrester Wave: Data Quality Solutions, Q1 2026. Forrester, Mar 2026. <https://www.forrester.com/blogs/the-forrester-wave-data-quality-solutions-q1-2026>
- [32] Data Governance blog category (SAP data control plane). Forrester, May 2026. <https://www.forrester.com/blogs/category/data-governance/>
- [33] Building a High-Performance Data and AI Organization. MIT Technology Review Insights, sponsored by Databricks, 2025. <https://www.databricks.com/sites/default/files/2025-10/mittr-databricks-data-and-ai-2025-report.pdf>
- [34] Donating the Model Context Protocol. Anthropic, Dec 2025. <https://www.anthropic.com/news/donating-the-model-context-protocol-and-establishing-of-the-agentic-ai-foundation>
- [35] Towards a Unified Digital Ecosystem. Mehmood et al., Electronics (MDPI), Dec 2025. <https://www.mdpi.com/2079-9292/14/24/4787>
- [36] Unified Data Platforms: Architecture, Benefits, and Implementation. Fivetran, Jun 2026. <https://www.fivetran.com/learn/unified-data-platform>
- [37] Vol. 01 State of Data Architecture; Vol. 05 State of AI Observability. Modern Data 101, 2026. <https://www.moderndata101.com>



540+ Respondents, 64 Countries, 29 Industries

[Access the report now](#)



Modern Data 101 Community



Subscribe 
moderndata101.com